arXiv:1511.02974v1 [math.OC] 10 Nov 2015

New Computational Guarantees for Solving Convex Optimization Problems with First Order Methods, via a Function Growth Condition Measure

Robert M. Freund* Haihao Lu†

November 11, 2015

Abstract

Motivated by recent work of Renegar [20], we present new computational methods and associated computational guarantees for solving convex optimization problems using first-order methods. Our problem of interest is the general convex optimization problem $f^* = \min_{x \in Q} f(x)$, where we presume knowledge of a strict lower bound $f_{\text{slb}} < f^*$. [Indeed, f_{slb} is naturally known when optimizing many loss functions in statistics and machine learning (least-squares, logistic loss, exponential loss, total variation loss, etc.) as well as in Renegar's transformed version of the standard conic optimization problem [20]; in all these cases one has $f_{\text{slb}} = 0 < f^*$.] We introduce a new functional measure called the growth constant G for $f(\cdot)$, that measures how quickly the level sets of $f(\cdot)$ grow relative to the function value, and that plays a fundamental role in the complexity analysis. When $f(\cdot)$ is non-smooth, we present new computational guarantees for the Subgradient Descent Method and for smoothing methods, that can improve existing computational guarantees in several ways, most notably when the initial iterate x^0 is far from the optimal solution set. When $f(\cdot)$ is smooth, we present a scheme for periodically restarting the Accelerated Gradient Method that can also improve existing computational guarantees when x^0 is far from the optimal solution set, and in the presence of added structure we present a scheme using parametrically increased smoothing that further improves the associated computational guarantees.

1 Problem Statement and Overview of Results

1.1 Problem Statement, Strict Lower Bound, and Function Growth Constant

Motivated by recent work of Renegar [20], we present new computational methods and associated computational guarantees for solving convex optimization problems using first-order methods. Our

*MIT Sloan School of Management, 77 Massachusetts Avenue, Cambridge, MA 02139 (mailto: rfreund@mit.edu). This author's research is supported by AFOSR Grant No. FA9550-15-1-0276 and the MIT-Belgium Université Catholique de Louvain Fund.

†MIT Department of Mathematics, 77 Massachusetts Avenue, Cambridge, MA 02139 (mailto: haihao@mit.edu).

problem of interest is the following optimization problem:

$$\begin{array}{rlll} P: & f^* := & \text{minimum}_x & f(x) \\ & & \text{s.t.} & x \in Q \ , \end{array} \tag{1}$$

where $Q \subseteq \mathbb{R}^n$ is a closed convex set and $f(\cdot) : Q \to \mathbb{R}$ is a convex function. Let the set of optimal solutions of (1) be denoted as $\text{Opt} := \{x \in Q : f(x) = f^*\}$. For $x \in Q$, let $\text{Dist}(x, \text{Opt})$ denote the distance from x to the set of optimal solutions, namely $\text{Dist}(x, \text{Opt}) := \min_y\{\|y - x\| : y \in \text{Opt}\}$.

Strict Lower Bound f_{slb} and Function Growth Constant G. Let f_{slb} be a known and given strict lower bound on the optimal value f^* of (1), namely $f_{\text{slb}} < f^*$. Such a known strict lower bound arises naturally when optimizing many loss functions in statistics and machine learning (least-squares loss, logistic loss, exponential loss, total variation loss, etc.) perhaps with the addition of a regularization term; in all these cases $f_{\text{slb}} = 0 < f^*$. A known strict lower bound also arises in Renegar's transformed version of the standard conic optimization problem [20].

Let $\varepsilon' > 0$ be given. Given the knowledge of the strict lower bound f_{slb}, it is natural to work with the notion of a relative measure of optimality. Let us define an ε'-relative solution of (1) to be a point $\hat{x}$ that satisfies:

$$\hat{x} \in Q \quad \text{and} \quad \frac{f(\hat{x}) - f^*}{f^* - f_{\text{slb}}} \ \leq \ \varepsilon' \ . \tag{2}$$

Note that (2) is a relative error measure, relative to the optimal bound gap $f^* - f_{\text{slb}}$. We focus on an ε'-relative solution rather than on an ε-absolute solution ($f(\hat{x}) \leq f^* + \varepsilon$), as the former seems more natural in the setting where a strict lower bound is part of the problem description. Indeed, consider the context of loss functions $f(\cdot)$ in statistics and machine learning where $f_{\text{slb}} = 0$, in which case an ε'-relative solution $\hat{x}$ corresponds to $\frac{f(\hat{x})}{f^*} \leq (1 + \varepsilon')$, and hence is a multiplicative measure of optimality tolerance.

Let G denote the smallest scalar $\bar{G}$ satisfying:

$$\text{Dist}(x, \text{Opt}) \leq \bar{G} \cdot (f(x) - f_{\text{slb}}) \ \text{ for all } x \in Q \ . \tag{3}$$

By its definition one sees that G measures how fast the distances from the optimal solution set Opt grow relative to the bound gap $f(x) - f_{\text{slb}}$. Therefore G is a measure of the growth rate of the level sets of $f(\cdot)$. We call G the "growth constant" of the function $f(\cdot)$ for the given strict lower bound f_{slb}. Note that an equivalent definition of G is given by:

$$G = \sup_{x \in Q} \left\{ \frac{\text{Dist}(x, \text{Opt})}{f(x) - f_{\text{slb}}} \right\} \ . \tag{4}$$

Unlike the strict lower bound f_{slb}, we do not assume that G is known, nor do we need any upper bounds on G in the computational methods we present and analyze.

We will see that the knowledge of the fixed strict lower bound f_{slb} and the concept of the function growth constant G leads to different versions of first-order methods with different computational

guarantees than the traditional analysis of first-order methods would dictate, and that these guarantees can dominate the classical guarantees in many cases but most notably when the initial iterate x^0 is far from the optimal solution. Roughly speaking, these computational guarantees grow at most as $\ln(f(x^0) - f^*)$ for non-smooth optimization, in contrast to standard analysis of such methods wherein the computational guarantees grows as $\text{Dist}(x^0, \text{Opt})^2$. For smooth optimization, the guarantees grow only as $\sqrt{f(x^0) - f^*}$ or $\ln(f(x^0) - f^*)$ — depending on the algorithm — in contrast to traditional analysis of such methods wherein the computational guarantees grow as $\text{Dist}(x^0, \text{Opt})$. Given the key role that G and f_{slb} play in the analysis to come, we now discuss several aspects of these definitions and use of these quantities.

In a departure from typical optimization approaches to lower bounds such as those arising from duality theory wherein one desires as tight a lower bound as possible, herein the lower bound f_{slb} is strict, namely $f_{\text{slb}} < f^*$, and it is fixed, i.e., it is not updated as part of a computational procedure. It is best to think of this lower bound as a *structural* lower bound that is easily connected to known properties of the function $f(\cdot)$. Such a strict lower bound on $f(\cdot)$ arises naturally in the settings of statistics and machine learning in the case of loss functions and/or regularization functions, see for example [10]. Consider when $f(\cdot)$ is the logistic loss function $f(x) = \frac{1}{m}\sum_{i=1}^m \ln\left(1 + e^{-A_i x}\right)$ or the exponential loss function $f(x) = \ln\left(\frac{1}{m}\sum_{i=1}^m e^{-A_i x}\right)$, perhaps with the addition of a regularization term $\lambda\|x\|_p^r$ for some $p \geq 1$, $r \geq 1$, and $\lambda \geq 0$. If the sample data is not strictly separable, which translated herein means that there is no x satisfying $Ax \geq 0$ unless $Ax = 0$, then it follows that $f^* > 0$ and so $f_{\text{slb}} := 0$ is a strict lower bound and is quite natural in this setting. Another example is regularized least-squares regression such as the LASSO and its cousins, wherein $f(\beta) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\beta\|^2 + \lambda\|\beta\|_p^r$; it follows that $f^* \geq 0$ and one can assert that $f^* > 0 =: f_{\text{slb}}$ under a variety of mild assumptions involving either λ or the data matrix $\mathbf{X}$. Other classes of examples for which $f_{\text{slb}} = 0$ is a strict lower bound on f^* include total variation (TV) loss functions which are used in image de-noising, as well as the broad class of minimum norm problems in general, under mild assumptions. Another class of problems for which there is a natural strict lower bound on f^* is the class of projectively transformed conic convex optimization problems under a particularly clever projective transformation, as developed by Renegar [20]; indeed it was this problem class and the results in [20] that gave rise to the line of research described herein.

We can interpret G as connected to a lower estimator of $f(\cdot)$: rearranging (3), we obtain:

$$f(x) \;\geq\; \bar{f}(x) := f_{\text{slb}} + G^{-1}\text{Dist}(x, \text{Opt}) \quad \text{for all } x \in Q\ . \tag{5}$$

Therefore the convex function $\bar{f}(x) = f_{\text{slb}} + G^{-1}\text{Dist}(x, \text{Opt})$ is a lower estimator of the function $f(\cdot)$ on Q. This interpretation is illustrated in Figure 1. As Figure 1 illustrates, the concept of the growth constant G is somewhat related to the notion of the modulus of weak sharp minima for (1), see Polyak [18] and Burke and Ferris [3]; this relationship is discussed further in Appendix A.1.

A natural question to ask is under what circumstances is the growth constant G finite? Roughly speaking, it holds that G is finite except when the objective function level sets are ill-behaved with regards to their recession cone. This is made precise in the following theorem, whose proof is given in Appendix A.2. For $\varepsilon > 0$, let $\text{Opt}_\varepsilon := \{x \in Q :\ f(x) \leq f^* + \varepsilon\}$ denote the ε-optimal level set of $f(\cdot)$ on Q, and let S denote the recession cone of Opt_ε. Note that S is the (common) recession cone of Opt_ε for all $\varepsilon \geq 0$.

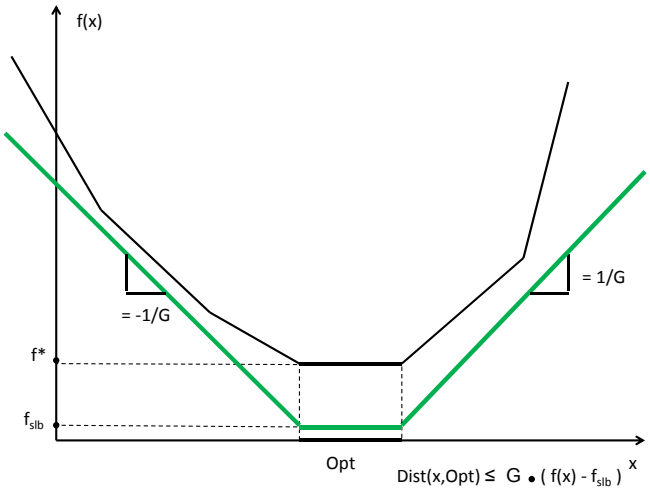

Figure 1: Illustration of G and f_{slb} for a function with multiple optimal solutions.

Theorem 1.1. *Suppose that for some $\varepsilon > 0$ there exists a bounded set E_ε for which $\text{Opt}_\varepsilon \subset E_\varepsilon + S$ where S is the recession cone of Opt_ε. Then for any given strict lower bound $f_{\text{slb}} < f^*$, the growth constant G is finite.* □

Let us briefly examine special cases of Theorem 1.1. In the case when Opt is bounded, then Theorem 1.1 implies G is finite. And in cases when $\text{Opt} = E + T$ where E is a bounded convex set and T is a subspace, then Theorem 1.1 implies that G is finite in these cases as well.

For an example wherein $G = \infty$, consider the function $f(x_1, x_2) := \frac{x_2^2}{x_1}$ on $Q := \{(x_1, x_2) : x_1 \geq 1\}$. Then $f(\cdot)$ is convex on Q, but the growth constant $G = \infty$ for any strict lower bound f_{slb}.

1.2 Overview of Results

We use the knowledge of the fixed strict lower bound f_{slb} and the concept of the function growth constant G to design and develop computational guarantees for new versions of first-order methods for solving the optimization problem (1). In Section 3 we present such methods when $f(\cdot)$ is non-smooth and Lipschitz continuous with Lipschitz constant M. In Theorem 3.1 we present a new step-size rule and analysis of Subgradient Descent that leads to an iteration complexity of $O(M^2G^2[\frac{\ln(H)}{\varepsilon'} + \frac{1}{(\varepsilon')^2}])$, where $H := \frac{f(x^0) - f_{\text{slb}}}{f^* - f_{\text{slb}}}$, which strictly improves the standard computational complexity bound for Subgradient Descent when x^0 is a "cold start," i.e., $\text{Dist}(x^0, \text{Opt})$ is large. In the special case when the optimal objective function value f^* is known, this iteration bound is further improved in Theorem 3.2. And when $f(\cdot)$ can be smoothed, we present further improved computational guarantees for a new method (Algorithm 3) that successively smooths and restarts the Accelerated Gradient Method, see Theorem 3.3 herein.

In Section 4 we present computational guarantees for new first-order methods when $f(\cdot)$ is smooth and has Lipschitz gradient with Lipschitz constant L. We present a new first-order method (Algorithm 4) based on periodically restarting the Accelerated Gradient Method, that leads to an iteration complexity of $O(G\sqrt{L}[\frac{\sqrt{f^*-f_{\rm slb}}}{\sqrt{\varepsilon'}} + \sqrt{f(x^0)-f_{\rm slb}}])$ (Theorem 4.1), which in many cases can improve the standard computational complexity bound for the Accelerated Gradient Method, most notably when x^0 is far from the optimal solution and ε' is small. And when $f(\cdot)$ has appropriate adjoint structure, we present another new method (Algorithm 5) that uses parametric increased smoothing and restarting of the Accelerated Gradient Method, and achieves a further modest improvement in the above computational guarantee (Theorem 4.2).

The paper is organized as follows. Section 2 contains a brief review of Subgradient Descent and an Accelerated Gradient Method. Section 3 contains first-order methods and computational guarantees when $f(\cdot)$ is non-smooth. Section 4 contains first-order methods and computational guarantees when $f(\cdot)$ is smooth.

Notation. Unless otherwise specified, the norm is the Euclidean (inner product) norm $\|x\| := \sqrt{x^T x}$. We occasionally refer to the ℓ_p norm of a vector v, which is denoted by $\|v\|_p$. For $Q \subset \mathbb{R}^n$, let $\Pi_Q(\cdot)$ denote the Euclidean projection operator onto Q, namely $\Pi_Q(x) := \arg\min_{y\in Q}\{\|y-x\|\}$. We define $\text{Dist}(x,S) := \min_y\{\|x-y\| : y \in S\}$. The set of optimal solutions of (1) is denoted by $\text{Opt} := \{x \in Q : f(x) = f^*\}$.

2 Review of Subgradient Descent and an Accelerated Gradient Method

We briefly review the Subgradient Descent Method and an Accelerated Gradient Method (as analyzed in Tseng [24]) for solving the convex optimization problem (1).

2.1 Subgradient Descent

Recall that g is a subgradient of $f(\cdot)$ at x if the following subgradient inequality holds:

$$f(y) \geq f(x) + g^T(y-x) \quad \text{for all } y \in Q\ . \tag{6}$$

Let $\partial f(x)$ denote the set of subgradients of $f(\cdot)$ at x. Here we assume that $f(\cdot)$ is Lipschitz continuous on a relatively open set Q^o containing Q, namely, there is a scalar M for which

$$|f(y)-f(x)| \leq M\|y-x\| \quad \text{for all } x,y \in Q^o\ . \tag{7}$$

It follows from (7) that for all for $x \in Q$ and $g \in \partial f(x)$ it holds that $\|g\| \leq M$.

Algorithm 1 presents the standard subgradient scheme. In this method x^k is the iterate at iteration k, the best objective value among the first k iterates is f_b^k, and the best iterate among the first k iterates is x_b^k.

Algorithm 1 Subgradient Method for Non-Smooth Optimization

Initialize. Initialize with $x^0 \in Q$, $f_b^0 \leftarrow f(x^0)$, $x_b^0 \leftarrow x^0$. $i \leftarrow 0$.

At iteration i :

1. **Compute Subgradient.** Compute $g_i \in \partial f(x^i)$.
2. **Determine Step-size.** Determine $\alpha_i \geq 0$.
3. **Perform Updates.** Compute $x^{i+1} \leftarrow \Pi_Q(x^i - \alpha_i g_i)$,

$$f_b^{i+1} \leftarrow \min\{f_b^i, f(x^{i+1})\} \ ,$$

$$x_b^{i+1} \leftarrow \underset{x \in \{x_b^i,\ x^{i+1}\}}{\arg\min} \{f(x)\} \ .$$

The following theorem summarize well-known computational guarantees associated with the subgradient descent method.

Theorem 2.1. (Convergence Bounds for Subgradient Descent [19, 14])
(i) Consider the subgradient descent method (Algorithm 1). Then for all $k \geq i \geq 0$, the following inequality holds:

$$f_b^k \ \leq \ f^* \ + \ \frac{\mathrm{Dist}(x^i, \mathrm{Opt})^2 + \sum_{l=i}^k \|g_l\|^2 \alpha_l^2}{2\sum_{l=i}^k \alpha_l} \ \leq \ f^* \ + \ \frac{\mathrm{Dist}(x^i, \mathrm{Opt})^2 + M^2 \sum_{l=i}^k \alpha_l^2}{2\sum_{l=i}^k \alpha_l} \ . \quad (8)$$

(ii) Suppose that f^* is known, and let the step-sizes for Algorithm 1 be $\alpha_i = (f(x^i) - f^*)/\|g_i\|^2$. Then for all $k \geq i \geq 0$, the following inequality holds:

$$f_b^k \ \leq \ f^* \ + \ \frac{M\mathrm{Dist}(x^i, \mathrm{Opt})}{\sqrt{k-i+1}} \ . \qquad \square \quad (9)$$

Suppose that we seek to bound the number of iterations N of the Subgradient Descent method required to compute an (absolute) ε-optimal solution of (1), which is a point $\hat{x} \in Q$ that satisfies $f(\hat{x}) \leq f^* + \varepsilon$. If $\varepsilon > 0$ is given, and the step-sizes are chosen as $\alpha_i = \varepsilon/\|g_i\|^2$, then it follows from Theorem 2.1 that $f(x_b^N) \leq f^* + \varepsilon$ for all

$$N \geq \bar{N} := \frac{M^2 \mathrm{Dist}(x^0, \mathrm{Opt})^2}{\varepsilon^2} - 1 \ . \quad (10)$$

If instead we know (or can bound from above) $\mathrm{Dist}(x^0, \mathrm{Opt})$, and the step-sizes are chosen as $\alpha_i = \mathrm{Dist}(x^0, \mathrm{Opt})/(\sqrt{N+1}\|g_i\|)$ where N satisfies (10), then it also follows from Theorem 2.1 that $f(x_b^N) \leq f^* + \varepsilon$. And if f^* is known, then the bound (10) is also sufficient to guarantee $f(x_b^N) \leq f^* + \varepsilon$ if the steps-sizes are chosen as in part (ii) of Theorem 2.1. Furthermore, it follows from [13] that the bound (10) cannot in general be improved in the black-box oracle model of computation. In this regard, we note that knowledge of the strict lower bound f_{slb}, which is used throughout this paper, falls outside of the assumptions of the black-box oracle model.

2.2 Accelerated Gradient Method for Smooth Optimization

Here we assume that $f(\cdot)$ is differentiable with Lipschitz continuous gradient on Q, namely, there is a scalar L for which

$$\|\nabla f(y) - \nabla f(x)\| \leq L\|y - x\| \quad \text{for all } x, y \in Q \ . \tag{11}$$

Algorithm 2 presents a standard Accelerated Gradient Method as in Tseng [24].

Algorithm 2 Accelerated Gradient Method

Initialize. Initialize with $x^0 \in Q$ and $z^0 := x^0$, and $i \leftarrow 0$. Define step-size parameters $\theta_i \in (0, 1]$ recursively by $\theta_0 := 1$ and θ_{i+1} satisfies $\frac{1}{\theta_{i+1}^2} - \frac{1}{\theta_{i+1}} = \frac{1}{\theta_i^2}$.

At iteration i :

1. **Perform Updates.** Define $y^i \leftarrow (1 - \theta_i)x^i + \theta_i z^i$, and compute $\nabla f(y^i)$,

$$z^{i+1} \leftarrow \arg\min_{x \in Q}\{f(y^i) + \nabla f(y^i)^T(x - z^i) + \tfrac{1}{2}\theta_i L(x - z^i)^T(x - z^i)\} \ ,$$

$$x^{i+1} \leftarrow (1 - \theta_i)x^i + \theta_i z^{i+1} \ .$$

For $\delta \geq f^*$ define the level set $S_\delta := \{x \in Q : f(x) \leq \delta\}$. For $x \in Q$, let $\text{Dist}(x, S_\delta)$ denote the distance from x to the level set S_δ, namely $\text{Dist}(x, S_\delta) := \min_y\{\|y - x\| : y \in S_\delta\}$. The following theorem is a computational guarantee for the Accelerated Gradient Method.

Theorem 2.2. (Convergence Bound for Accelerated Gradient Method [24]) *Consider the Accelerated Gradient Method (Algorithm 2). Let $\delta \geq f^*$ and $S_\delta := \{x \in Q : f(x) \leq \delta\}$. Then for all $k \geq 0$, the following inequality holds:*

$$f(x^k) - \delta \quad \leq \quad \frac{2L\text{Dist}(x^0, S_\delta)^2}{(k+1)^2} \ . \qquad \square \tag{12}$$

3 Computational Guarantees when $f(\cdot)$ is Non-Smooth

Let $\varepsilon' > 0$ be given. We aspire to compute an ε'-relative solution of (1), which recall from (2) is a point $\hat{x} \in Q$ satisfying: $\frac{f(\hat{x}) - f^*}{f^* - f_{\text{slb}}} \leq \varepsilon'$. In this section we present three new computational guarantees for first-order methods applied to computing a ε'-relative solution of problem (1) that are based on the strict lower bound f_{slb} and growth constant G. The first two guarantees are for the Subgradient Descent Method, and they differ according to the information required to compute step-sizes for the method. The first guarantee is for the case when only the strict lower bound f_{slb} is known. The second guarantee is for the case when the optimal value f^* is known. The third guarantee is for the case when the function $f(\cdot)$ can be smoothed and then solved using the Accelerated Gradient Method.

3.1 Subgradient Descent with New Step-Size Rules

We first consider solving (1) using the Subgradient Descent Method when the strict lower bound $f_{\rm slb}$ is known. We have:

Theorem 3.1. (Complexity Bound for Subgradient Descent when $f_{\rm slb}$ is known) *Let the relative accuracy ε' be given, and let the step-sizes for the Subgradient Descent Method applied to solve problem* (1) *be chosen as:*

$$\alpha_i := \left(\frac{f_b^i - f_{\rm slb}}{\sqrt[3]{e}\|g_i\|^2}\right)\left(\frac{\varepsilon'}{1+\varepsilon'}\right) \ ,$$

and suppose that

$$N \ \geq \ M^2G^2\left(16\left(\frac{1+\varepsilon'}{\varepsilon'}\right)\ln\left(1+\frac{f(x^0)-f^*}{f^*-f_{\rm slb}}\right) \ + \ 11\left(\frac{1+\varepsilon'}{\varepsilon'}\right)^2\right) \ .$$

Then:

$$\frac{f(x_b^N)-f^*}{f^*-f_{\rm slb}} \ \leq \ \varepsilon' \ . \qquad \square$$

This step-size rule and associated guarantee generalizes Theorem 3.8 of Renegar [20], which therein pertains to a specific transformed conic optimization problem. Here the step-sizes are chosen in a "smooth" manner, as opposed to the rule given in Algorithm 2 of [20].

Let us compare the computational guarantee for Subgradient Descent with the step-size rule based on the strict lower bound $f_{\rm slb}$ (Theorem 3.1), with the standard guarantee for Subgradient Descent given by (10). From (2) we have that an ε'-relative solution corresponds to an absolute $\varepsilon = \varepsilon' \cdot (f^* - f_{\rm slb})$ solution of (1). Let us presume that ε' is small, whereby $\frac{1+\varepsilon'}{\varepsilon'} \approx \frac{1}{\varepsilon'}$. Define $H := \frac{f(x^0)-f_{\rm slb}}{f^*-f_{\rm slb}}$. Then the standard computational guarantee (10) can be written as:

$$\frac{M^2{\rm Dist}(x^0,{\rm Opt})^2}{\varepsilon'^2\left(f^*-f_{\rm slb}\right)^2} \ ,$$

and the new computational guarantee using Theorem 3.1 is then written as:

$$M^2G^2\left(16\left(\frac{1}{\varepsilon'}\right)\ln(H) + 11\left(\frac{1}{\varepsilon'^2}\right)\right) \ .$$

Let us now consider cases wherein ${\rm Dist}(x^0,{\rm Opt}) > G\cdot(f^*-f_{\rm slb})$. Since G and $f^*-f_{\rm slb}$ are themselves fixed for a given problem instance, this means that ${\rm Dist}(x^0,{\rm Opt})$ is large, i.e., x^0 is far from the set of optimal solutions. Let $\bar{C}$ denote the ratio of ${\rm Dist}(x^0,{\rm Opt})$ to $G\cdot(f^*-f_{\rm slb})$, namely

$$\bar{C} := \frac{{\rm Dist}(x^0,{\rm Opt})}{G(f^*-f_{\rm slb})} > 1 \ . \tag{13}$$

Note that $f(x^0) \leq f^* + M{\rm Dist}(x^0,{\rm Opt})$, and therefore

$$H = 1 + \frac{f(x^0)-f^*}{f^*-f_{\rm slb}} \leq 1 + \frac{M{\rm Dist}(x^0,{\rm Opt})}{f^*-f_{\rm slb}} = 1 + MG\bar{C} \ .$$

Therefore the ratio of the new guarantee in Theorem 3.1 to the standard guarantee (10) is at most

$$\frac{\text{Guarantee of Theorem 3.1}}{\text{Standard Guarantee (10)}} \;\leq\; \frac{16\varepsilon' \ln\left(1+MG\bar{C}\right)}{\bar{C}^2} \;+\; \frac{11}{\bar{C}^2}\;. \tag{14}$$

Notice from (13) that when $\text{Dist}(x^0, \text{Opt})$ is sufficiently large that the ratio $\bar{C}$ is very large and the right-hand side of (14) can be made arbitrarily small, thereby showing that in these cases the computational guarantee in Theorem 3.1 can be made arbitrarily better than the standard guarantee (10) for Subgradient Descent. In fact, even if $\bar{C}$ is not particularly large (say, $\bar{C} = 10$), then for small values of ε' it will still hold that the computational guarantee in Theorem 3.1 will be better than the standard guarantee (10).

We will prove Theorem 3.1 by proving a series of five propositions that, when taken together, will provide the necessary demonstration of the theorem. The reader familiar with [20] will notice many resemblances between the proof below and the proof of Theorem 3.8 of [20], especially so in that Theorem 3.1 is a generalization thereof to a broad class of convex functions. For notational convenience, we will work with three constants B, F, and $\hat{R}$, whose specific values we will set later on. These constants will satisfy the following conditions:

$$B \in (0,1) \;\;,\;\; F > \frac{1}{2B} \;\;,\text{ and }\;\; \hat{R} > \frac{1-\frac{1}{2F}}{B-\frac{1}{2F}}\;. \tag{15}$$

Given ε', it will be notationally convenient to define $\varepsilon := \frac{\varepsilon'}{1+\varepsilon'}$. We consider setting the step-sizes for Subgradient Descent to be

$$\alpha_i := \left(\frac{f_b^i - f_{\text{slb}}}{F\|g_i\|^2}\right)\varepsilon\;,$$

and hence the step-size rule in Theorem 3.1 uses $F = \sqrt[3]{e}$ where $e = 2.718...$ is the base of the natural logarithm.

Define $K_0 := 0$, and since $B < 1$, for all i define K_{i+1} inductively as the smallest iteration index of Subgradient Descent for which $f(x^{K_{i+1}}) - f_{\text{slb}} \leq B(f(x^{K_i}) - f_{\text{slb}})$.

Proposition 3.1. *For $j = 0, 1, \ldots, K_{i+1} - K_i - 1$, it holds that*

$$f(x_b^{K_i+j}) - f_{\text{slb}} \leq f^* - f_{\text{slb}} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(j+1)} + \frac{\varepsilon}{2F}\right](f(x^{K_i}) - f_{\text{slb}})\;.$$

Proof: It follows from Theorem 2.1 that

$$\begin{aligned} f(x_b^{K_i+j}) - f_{\text{slb}} \;&\leq\; f^* - f_{\text{slb}} + \frac{\text{Dist}(x^{K_i}, \text{Opt})^2}{2\sum_{l=K_i}^{K_i+j}\alpha_l} + \frac{\sum_{l=K_i}^{K_i+j}\|g_l\|^2\alpha_l^2}{2\sum_{l=K_i}^{K_i+j}\alpha_l} \\ &\leq\; f^* - f_{\text{slb}} + \frac{G^2(f_b^{K_i} - f_{\text{slb}})^2}{2\sum_{l=K_i}^{K_i+j}\alpha_l} + \frac{\sum_{l=K_i}^{K_i+j}\|g_l\|^2\alpha_l^2}{2\sum_{l=K_i}^{K_i+j}\alpha_l}, \end{aligned} \tag{16}$$

where the second inequality uses the definition of G directly and the fact that $f_b^{K_i} = f(x^{K_i})$. Since the sequence $\{f_b^i\}$ is decreasing in i, and $\|g_i\| \leq M$, it holds that

$$\sum_{l=K_i}^{K_i+j}\alpha_l \geq (j+1)\min_{l=K_i,\ldots,K_i+j}\{\alpha_l\} \geq (j+1)\frac{(f_b^{K_i+j} - f_{\text{slb}})\varepsilon}{FM^2} > \frac{(j+1)\varepsilon B(f_b^{K_i} - f_{\text{slb}})}{FM^2}\;,$$

where the strict inequality follows since $K_i+j < K_{i+1}$ and hence $f(x^{K_i+j})-f_{\rm slb} > B(f(x^{K_i})-f_{\rm slb})$. Also since the sequence $\{f_b^i\}$ is decreasing in i, it holds that

$$\frac{\sum_{l=K_i}^{K_i+j}\|g_l\|^2\alpha_l^2}{\sum_{l=K_i}^{K_i+j}\alpha_l} \leq \max_{l=K_i,\ldots,K_i+j}\{\|g_l\|^2\alpha_l\} = \max_{l=K_i,\ldots,K_i+j}\left\{\frac{(f_b^l-f_{\rm slb})\varepsilon}{F}\right\} = \frac{\varepsilon(f_b^{K_i}-f_{\rm slb})}{F} \ .$$

Substituting these inequalities into (16) yields the result. □

Proposition 3.2. *Define:*

$$U := \left\lfloor \frac{B^{-1}FM^2G^2}{2\varepsilon^2\left[B-\frac{1}{2F}\right]} \right\rfloor \ .$$

Then either $\frac{f(x_b^{K_i+U})-f^*}{f^*-f_{\rm slb}} \leq \varepsilon'$ *or* $K_{i+1}-K_i \leq U$.

Proof: Suppose that $\frac{f(x_b^{K_i+U})-f^*}{f^*-f_{\rm slb}} > \varepsilon'$. Then it holds that:

$$\varepsilon' < \frac{f(x_b^{K_i+U})-f_{\rm slb}-(f^*-f_{\rm slb})}{f^*-f_{\rm slb}} = \frac{f(x_b^{K_i+U})-f_{\rm slb}}{f^*-f_{\rm slb}} - 1 \ ,$$

whereby

$$\frac{f^*-f_{\rm slb}}{f(x_b^{K_i+U})-f_{\rm slb}} < \frac{1}{1+\varepsilon'} = 1-\varepsilon \ . \tag{17}$$

Let us also suppose that $K_{i+1}-K_i > U$; we will derive a contradiction from this supposition, which will furnish the proof. Let $j = U < K_{i+1}-K_i$. Then it follows from Proposition 3.1 that

$$f(x_b^{K_i+U})-f_{\rm slb} \leq f^*-f_{\rm slb} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(U+1)}+\frac{\varepsilon}{2F}\right](f(x^{K_i})-f_{\rm slb}) \ ,$$

and dividing by $f(x_b^{K_i+U})-f_{\rm slb}$ yields:

$$\begin{aligned} 1 &\leq \frac{f^*-f_{\rm slb}}{f(x_b^{K_i+U})-f_{\rm slb}} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(U+1)}+\frac{\varepsilon}{2F}\right]\frac{f(x^{K_i})-f_{\rm slb}}{f(x_b^{K_i+U})-f_{\rm slb}} \\ &< 1-\varepsilon+\left[\frac{B^{-1}G^2M^2F}{2\varepsilon(U+1)}+\frac{\varepsilon}{2F}\right]B^{-1} \ , \end{aligned} \tag{18}$$

where the strict inequality uses (17) and the fact that $K_i+U < K_{i+1}$ which implies that $B(f(x^{K_i})-f_{\rm slb}) < f(x_b^{K_i+U}f)-f_{\rm slb}$. Finally, (18) rearranges to yield:

$$U+1 < \frac{B^{-1}FM^2G^2}{2\varepsilon^2\left[B-\frac{1}{2F}\right]} < U+1 \ ,$$

which is a contradiction. □

For the next proposition, recall the constant $\hat{R}$ that must satisfy (15).

Proposition 3.3. *Suppose that* $\frac{f(x^{K_i})-f_{\rm slb}}{f^*-f_{\rm slb}} > \hat{R}$. *Define:*

$$V := \left\lfloor \frac{B^{-1}FM^2G^2}{2\varepsilon\left[B - \frac{1}{2F} - \frac{1}{\hat{R}}\left(1-\frac{1}{2F}\right)\right]} \right\rfloor \ .$$

Then either $\frac{f(x_b^{K_i+V})-f^*}{f^*-f_{\rm slb}} \leq \varepsilon'$ *or* $K_{i+1} - K_i \leq V$.

Proof: Suppose that $\frac{f(x_b^{K_i+V})-f^*}{f^*-f_{\rm slb}} > \varepsilon'$. Then it also holds that $\frac{f(x^{K_i})-f^*}{f^*-f_{\rm slb}} > \varepsilon'$, whereby

$$\frac{f(x^{K_i}) - f_{\rm slb}}{f^* - f_{\rm slb}} > \max\{\hat{R}, 1+\varepsilon'\} \ ,$$

and hence

$$\frac{f^* - f_{\rm slb}}{f(x^{K_i}) - f_{\rm slb}} < \min\left\{\frac{1}{\hat{R}} \ , \ \frac{1}{1+\varepsilon'}\right\} = \min\left\{\frac{1}{\hat{R}} \ , \ 1-\varepsilon\right\} \ . \tag{19}$$

Let us also suppose that $K_{i+1} - K_i > V$; we will derive a contradiction from this supposition, which will furnish the proof. Let $j = V < K_{i+1} - K_i$. Then it follows from Proposition 3.1 that

$$f(x_b^{K_i+V}) - f_{\rm slb} \leq f^* - f_{\rm slb} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(V+1)} + \frac{\varepsilon}{2F}\right](f(x^{K_i}) - f_{\rm slb}) \ ,$$

and dividing by $f(x^{K_i}) - f_{\rm slb}$ yields:

$$\begin{aligned} B < \frac{f(x_b^{K_i+V}) - f_{\rm slb}}{f(x^{K_i}) - f_{\rm slb}} &\leq \frac{f^* - f_{\rm slb}}{f(x^{K_i}) - f_{\rm slb}} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(V+1)} + \frac{\varepsilon}{2F}\right] \\ &< \min\left\{\frac{1}{\hat{R}} \ , \ 1-\varepsilon\right\} + \left[\frac{B^{-1}G^2M^2F}{2\varepsilon(V+1)} + \frac{\varepsilon}{2F}\right] \ , \end{aligned} \tag{20}$$

where the first strict inequality uses the fact that $K_i + V < K_{i+1}$ which implies that $B(f(x^{K_i}) - f_{\rm slb}) < f(x_b^{K_i+V}) - f_{\rm slb}$, and the second strict inequality uses (19). This then rearranges to:

$$V + 1 < \frac{B^{-1}FM^2G^2}{2\varepsilon\left[B - \min\left\{\frac{1}{\hat{R}} \ , \ 1-\varepsilon\right\} - \frac{\varepsilon}{2F}\right]} \ .$$

Now notice that

$$\min\left\{\frac{1}{\hat{R}} \ , \ 1-\varepsilon\right\} + \frac{\varepsilon}{2F} \leq \left(1 - \frac{1}{2F}\right)\left(\frac{1}{\hat{R}}\right) + \frac{1}{2F}(1-\varepsilon) + \frac{\varepsilon}{2F} = \frac{1}{2F} + \frac{1}{\hat{R}}\left(1 - \frac{1}{2F}\right) \ ,$$

whereby

$$V + 1 < \frac{B^{-1}FM^2G^2}{2\varepsilon\left[B - \frac{1}{2F} - \frac{1}{\hat{R}}\left(1-\frac{1}{2F}\right)\right]} < V + 1 \ ,$$

which is a contradiction. □

Proposition 3.4. *Let m denote the smallest value of i for which*

$$\frac{f(x^{K_i}) - f_{\rm slb}}{f^* - f_{\rm slb}} \leq \hat{R} \ .$$

Then

$$m \leq \left\lceil \frac{\ln\left(\frac{f(x^0)-f_{\rm slb}}{f^*-f_{\rm slb}}\right) - \ln(\hat{R})}{\ln(1/B)} \right\rceil \ .$$

Proof: If follows using induction that $f(x^{K_i}) - f_{\rm slb} \leq B^i(f(x^0) - f_{\rm slb})$. Therefore $\hat{R}(f^* - f_{\rm slb}) < f(x^{K_{m-1}}) - f_{\rm slb} \leq B^{m-1}(f(x^0) - f_{\rm slb})$, and taking logarithms yields

$$m - 1 < \frac{\ln\left(\frac{f(x^0) - f_{\rm slb}}{f^* - f_{\rm slb}}\right) - \ln(\hat{R})}{\ln(1/B)} \ ,$$

from which the result follows. □

Proposition 3.5. *Let p denote the largest value of i for which K_i exists (is finite). Then*

$$p \leq m + \left\lfloor \frac{\ln(\hat{R})}{\ln(1/B)} \right\rfloor \ .$$

Proof: Clearly $p \geq m$. We have $f^* - f_{\rm slb} \leq f(x^{K_p}) - f_{\rm slb} \leq B^{p-m}(f(x^{K_m}) - f_{\rm slb}) \leq B^{p-m}\hat{R}(f^* - f_{\rm slb})$, where we have used the definition of m in Proposition 3.4. Taking logarithms yields

$$p - m \leq \frac{\ln(\hat{R})}{\ln(1/B)} \ ,$$

from which the result follows. □

We now use these propositions to prove Theorem 3.1.

Proof of Theorem 3.1: Let us choose the following values for constants: $B = 1/\sqrt[3]{e}$, $F = \sqrt[3]{e}$, and $\hat{R} = 3.7$. It then follows from Propositions 3.4 and 3.5 that $m < 3\ln\left(\frac{f(x^0)-f_{\rm slb}}{f^*-f_{\rm slb}}\right)$ and $p \leq m+3$. Define

$$\begin{aligned} N_1 := \textstyle\sum_{i=1}^m V &= m \left\lfloor \frac{B^{-1}FM^2G^2}{2\varepsilon\left[B - \frac{1}{2F} - \frac{1}{\hat{R}}\left(1 - \frac{1}{2F}\right)\right]} \right\rfloor \\ &< 3\ln\left(\frac{f(x^0) - f_{\rm slb}}{f^* - f_{\rm slb}}\right) \times \frac{B^{-1}FM^2G^2}{2\varepsilon\left[B - \frac{1}{2F} - \frac{1}{\hat{R}}\left(1 - \frac{1}{2F}\right)\right]} \\ &< \frac{16M^2G^2\ln\left(\frac{f(x^0) - f_{\rm slb}}{f^* - f_{\rm slb}}\right)}{\varepsilon} \ , \end{aligned}$$

where the first inequality is from the bound on m in Proposition 3.4, and the second inequality is obtained by substituting in the values of the constants B, F, and $\hat{R}$. It follows from Proposition 3.3, and using induction, that up to the first N_1 iterations that either $\frac{f_b^{N_1}-f^*}{f^*-f_{\rm slb}} \leq \varepsilon'$ or $K_m = \sum_{i=1}^{m} K_i - K_{i-1} \leq \sum_{i=1}^{m} V = N_1$. If the former is true, then the theorem is proved since $N \geq N_1$. Let us therefore suppose that the latter is true, and define:

$$N_2 := \sum_{i=m+1}^{p+1} U = (p-m+1)\left\lfloor \frac{B^{-1}FM^2G^2}{2\varepsilon^2\left(B-\frac{1}{2F}\right)} \right\rfloor \leq \frac{4B^{-1}FM^2G^2}{2\varepsilon^2\left(B-\frac{1}{2F}\right)} < \frac{11M^2G^2}{\varepsilon^2} \ ,$$

where the first inequality is from the bound on p in Proposition 3.5, and the second inequality is obtained by substituting in the values of the constants B and F. It follows from Proposition 3.2, and using induction, that by the time the algorithm has computed iteration $K_m + N_2$ that $\frac{f_b^{K_m+N_2}-f^*}{f^*-f_{\rm slb}} \leq \varepsilon'$. Therefore since $N \geq N_1 + N_2 \geq K_m + N_2$ the result follows. □

We next consider solving (1) using the Subgradient Descent Method when the optimal value f^* is known. The following theorem presents a computational guarantee in this setting.

Theorem 3.2. (Complexity Bound for Subgradient Descent when f^* is known) *Let the step-sizes for the Subgradient Descent Method applied to solve problem* (1) *be chosen as:*

$$\alpha_i := \frac{f(x^i)-f^*}{\|g_i\|^2} \ ,$$

and suppose that

$$N \ \geq \ 2M^2G^2\left(1 + 2.9\ln\left(\frac{f(x^0)-f^*}{f^*-f_{\rm slb}}\right) + 2.9\ln\left(\frac{1}{\varepsilon'}\right) + 6.8\left(\frac{1}{\varepsilon'}\right) + 2\left(\frac{1}{\varepsilon'}\right)^2\right) \ .$$

Then:

$$\frac{f(x_b^N)-f^*}{f^*-f_{\rm slb}} \ \leq \ \varepsilon' \ . \qquad \square$$

This step-size rule and associated guarantee is an almost-exact generalization of Theorem 3.7 of Renegar [20], which therein pertains to a specific transformed conic optimization problem. The proof of this theorem follows the logic for the proof of Theorem 3.7 of [20] in many respects as well.

Let us briefly compare the computational guarantees of Theorems 3.2 and 3.1. For ε' small, the dominant terms in the iteration bound in Theorem 3.2 are the second and fifth terms in the parentheses, and after removing constants and the multiplicative factor M^2G^2, we see that the key difference between the two bounds is that the bound in Theorem 3.2 has the term $\ln\left(\frac{f(x^0)-f^*}{f^*-f_{\rm slb}}\right)$, whereas the corresponding term in the bound in Theorem 3.1 involves $\left(\frac{1}{\varepsilon'}\right)\ln\left(\frac{f(x^0)-f_{\rm slb}}{f^*-f_{\rm slb}}\right)$. Therefore for small values of ε' and/or for large values of $f(x^0)$ the results of Theorem 3.2 improve on those of Theorem 3.1.

Proof of Theorem 3.2: Let $B \in (0,1)$ be a given fractional quantity. Define $K_0 := 0$, and since $B < 1$, for all i define K_{i+1} inductively as the smallest iteration index of Subgradient Descent for which $f(x^{K_{i+1}}) - f^* \leq B(f(x^{K_i}) - f^*)$. It follows from part (ii) of Theorem 2.1 that:

$$f(x_b^{K_{i+1}-1}) - f^* \leq \frac{M\text{Dist}(x^{K_i}, \text{Opt})}{\sqrt{K_{i+1} - 1 - K_i + 1}} \ .$$

This last inequality can be rearranged to yield:

$$K_{i+1} - K_i \leq \frac{M^2\text{Dist}(x^{K_i}, \text{Opt})^2}{\left(f(x_b^{K_{i+1}-1}) - f^*\right)^2} < \frac{B^{-2}M^2G^2\left(f(x^{K_i}) - f_{\text{slb}}\right)^2}{\left(f(x^{K_i}) - f^*\right)^2} = B^{-2}M^2G^2\left(1 + \frac{f^* - f_{\text{slb}}}{f(x^{K_i}) - f^*}\right)^2 \tag{21}$$

where the second inequality uses the definition of the growth constant G as well as the fact that $f(x^{K_{i+1}-1}) - f^* > B(f(x^{K_i}) - f^*)$.

Now let i' be the first sub-index that satisfies $\frac{f(x^{K_{i'}}) - f^*}{f^* - f_{\text{slb}}} \leq \varepsilon'$. Then it holds for any $i < i'$ that

$$\varepsilon'(f^* - f_{\text{slb}}) < f(x^{K_{i'-1}}) - f^* \leq B^{i'-1-i}(f(x^{K_i}) - f^*) \ . \tag{22}$$

Using $i = 0$ in (22) and taking logarithms yields:

$$i' < 1 + \frac{\ln\left(\frac{1}{\varepsilon'}\right) + \ln\left(\frac{f(x^0) - f^*}{f^* - f_{\text{slb}}}\right)}{\ln(B^{-1})} \ . \tag{23}$$

Putting all of this together we obtain:

$$\begin{aligned}
K_{i'} &= \sum_{i=0}^{i'-1}(K_{i+1} - K_i) \\
&\leq B^{-2}M^2G^2\sum_{i=0}^{i'-1}\left(1 + \frac{f^* - f_{\text{slb}}}{f(x^{K_i}) - f^*}\right)^2 \\
&\leq B^{-2}M^2G^2\sum_{i=0}^{i'-1}\left(1 + \frac{1}{\varepsilon'}B^{i'-1-i}\right)^2 \\
&= B^{-2}M^2G^2\sum_{j=0}^{i'-1}\left(1 + \left(\frac{2}{\varepsilon'}\right)B^j + \left(\frac{1}{\varepsilon'}\right)^2(B^2)^j\right) \\
&\leq B^{-2}M^2G^2\left(i' + \left(\frac{2}{\varepsilon'}\right)\frac{1}{1-B} + \left(\frac{1}{\varepsilon'}\right)^2\frac{1}{1-B^2}\right) \\
&\leq B^{-2}M^2G^2\left(1 + \frac{\ln\left(\frac{1}{\varepsilon'}\right) + \ln\left(\frac{f(x^0) - f^*}{f^* - f_{\text{slb}}}\right)}{\ln(B^{-1})} + \left(\frac{2}{\varepsilon'}\right)\frac{1}{1-B} + \left(\frac{1}{\varepsilon'}\right)^2\frac{1}{1-B^2}\right) \ ,
\end{aligned}$$

where the first inequality is from (21), the second inequality uses (22), the third inequality replaces the finite series with an infinite series, and the fourth inequality uses (23). Finally, using the value of $B = 1/\sqrt{2}$ and substituting into the above yields the result. $\square$

We remark that one obtains the precise constants of Theorem 3.7 of [20] by using $B = 1/2$. Choosing B to optimize the absolute constant of the $(1/\varepsilon')^2$ term yields $B = 1/\sqrt{2}$ and the absolute constants as presented in the statement of the threorem. Choosing B to optimize the absolute constant of the $\ln\left(\frac{f(x^0)-f^*}{f^*-f_{\text{slb}}}\right)$ term would yield $B = 1/\sqrt{e}$ with the coefficient of 2 in the $\ln(\cdot)$ terms.

3.2 Non-Smooth Optimization using a New Smooth Approximations Method

As first proposed by Nesterov [15], there are many practical settings wherein one can approximate the non-smooth convex function $f(\cdot)$ by a smooth convex function $f_\mu(\cdot)$, where the sense of the approximation depends on the parameter μ. If the smooth approximation $f_\mu(\cdot)$ is computationally easy to work with, one can then use the Accelerated Gradient Method (Algorithm 2) to approximately optimize $f_\mu(\cdot)$ (thereby also approximately optimizing $f(\cdot)$) on the feasible set Q. There are a variety of techniques that can be used to construct a parametric family of smooth functions $f_\mu(\cdot)$ depending on the known structure of $f(\cdot)$ and Q, see [15] as well as [16] and Beck and Teboulle [1] among others. For our purposes herein, we will suppose that there is a smoothing technique with the following two properties:

(i) there is a known constant $\bar{D} > 0$ such that for any given $\mu > 0$ we can construct a smooth convex function $f_\mu(\cdot) : Q \to \mathbb{R}$ which satisfies:

$$f(x) \leq f_\mu(x) \leq f(x) + \bar{D}\mu \text{ for all } x \in Q \text{ , and} \tag{24}$$

(ii) $f_\mu(\cdot)$ has Lipschitz continuous gradient on Q with Lipschitz constant $L_\mu \leq A/\mu$ for some known constant A.

These properties can be used to design an implementation of the Accelerated Gradient Method (Algorithm 2) applied to $f_\mu(\cdot)$, that can be used to compute an absolute ε-optimal solution of the original optimization problem (1). Using the scheme developed in [15] in conjunction with the Accelerated Gradient Method (Algorithm 2) yields an iteration complexity bound of

$$\check{N} := \left\lceil \frac{\sqrt{8A}\sqrt{\bar{D}}\text{Dist}(x^0, \text{Opt})}{\varepsilon} - 1 \right\rceil \tag{25}$$

for a suitably designed version of the basic method [15].

Herein we develop a variant of the basic smoothing method to solve the optimization problem (1) that yields a new computational guarantee that can improve on (25) in many cases. Algorithm 3 presents our method for computing an ε'-relative solution of the optimization problem (1) for non-smooth objective function $f(\cdot)$ based on successive smooth approximations and re-starting of the Accelerated Gradient Method (Algorithm 2). In the description of Algorithm 3 the notation "$x_{i,j} \leftarrow \text{AGM}(f_{\mu_i}(\cdot),\ x_{i,0},\ j)$" denotes assigning to $x_{i,j}$ the j^{th} iterate of the Accelerated Gradient Method applied to the optimization problem (1) with objective function $f_{\mu_i}(\cdot)$ using the initial point $x_{i,0} \in Q$.

Algorithm 3 Smooth Approximations Method

Initialize. Initialize with $x^0 \in Q$ and $\varepsilon' > 0$.
Set $x_{1,0} \leftarrow x^0$, $i \leftarrow 1$.

At outer iteration i :

1. **Set smoothing parameter.** $\mu_i \leftarrow \dfrac{\varepsilon' \cdot (f(x_{i,0}) - f_{\rm slb})}{5\bar{D}}$.
2. **Initialize inner iteration.** $j \leftarrow 0$
3. **Run inner iterations.** At inner iteration j:

If $\dfrac{f(x_{i,j}) - f_{\rm slb}}{f(x_{i,0}) - f_{\rm slb}} \geq 0.8$, then

$x_{i,j+1} \leftarrow \mathrm{AGM}(f_{\mu_i}(\cdot),\ x_{i,0},\ j+1)$,

$j \leftarrow j+1$, and Goto step 3.

Else $K_i \leftarrow j$, $x_{i+1,0} \leftarrow x_{i,j}$, $i \leftarrow i+1$, and Goto step 1.

At the $i^{\rm th}$ outer iteration of Algorithm 3, the algorithm sets the smoothing parameter μ_i in Step (1.) and then starts the Accelerated Gradient Method at the point $x_{i,0}$, and uses the double indexing notation of $x_{i,j}$ to denote iteration j of the Accelerated Gradient Method initialized at the point $x_{i,0}$ for optimizing $f_{\mu_i}(\cdot)$ on Q. Notice that the smoothing parameter μ_i decreases over the course of the outer iterations, as it makes more sense to set μ_i higher at first and then decrease it as the solution is approached. The outer iteration i runs until the ratio test in Step (3.) fails, at which point the current point $x_{i,j}$ becomes the starting point of the next outer iteration, namely $x_{i+1,0} \leftarrow x_{i,j}$. The counter K_i records the number of inner iterations computed in outer iteration i. Therefore the total number of inner iterations of Algorithm 3 after completing i outer iterations is $K_1 + \ldots + K_i$.

We have the following computational guarantee associated with Algorithm 3.

Theorem 3.3. (Complexity Bound for Smooth Approximations Method) *Let $x^0 \in Q$ be the initial point and let the relative accuracy $\varepsilon' \in (0,1]$ be given, and let x^N denote the iterate value of the Smooth Approximations Method (Algorithm 3) after a total of N inner iterations. If*

$$N \ \geq \ G\sqrt{A}\sqrt{\bar{D}} \left(32 \left[\frac{\ln\left(1 + \frac{f(x^0)-f^*}{f^*-f_{\rm slb}}\right)}{\sqrt{\varepsilon'}} \right] + 44 \left[\frac{1}{\varepsilon'} \right] \right) \ ,$$

then

$$\frac{f(x^N) - f^*}{f^* - f_{\rm slb}} \ \leq \ \varepsilon' \ . \qquad \square$$

Let us compare the computational guarantee of Theorem 3.3 with the standard guarantee for smoothing given in (25). Our discussion will be very similar to the discussion following Theorem 3.1. From (2) we have that an ε'-relative solution corresponds to an absolute $\varepsilon = \varepsilon' \cdot (f^* - f_{\rm slb})$ solution of (1). We define $H := \frac{f(x^0)-f_{\rm slb}}{f^*-f_{\rm slb}}$. Then the computational guarantee (25) can be written

as:

$$\frac{\sqrt{8A\bar{D}}\text{Dist}(x^0, \text{Opt})}{\varepsilon'\left(f^* - f_{\text{slb}}\right)} \ ,$$

and the new computational guarantee using Theorem 3.3 is then written as:

$$G\sqrt{A}\sqrt{\bar{D}}\left(32\left[\frac{\ln\left(H\right)}{\sqrt{\varepsilon'}}\right] + 44\left[\frac{1}{\varepsilon'}\right]\right) \ .$$

Let us now consider cases wherein $\text{Dist}(x^0, \text{Opt}) > G \cdot (f^* - f_{\text{slb}})$. Since G and $f^* - f_{\text{slb}}$ are themselves fixed for a given problem instance, this means that $\text{Dist}(x^0, \text{Opt})$ is large, i.e., x^0 is far from the set of optimal solutions. Let $\bar{C}$ denote the ratio of $\text{Dist}(x^0, \text{Opt})$ to $G \cdot (f^* - f_{\text{slb}})$, namely

$$\bar{C} := \frac{\text{Dist}(x^0, \text{Opt})}{G(f^* - f_{\text{slb}})} > 1 \ . \tag{26}$$

Note that $f(x^0) \leq f^* + M\text{Dist}(x^0, \text{Opt})$ where M is the Lipschitz constant for $f(\cdot)$, and therefore

$$H = 1 + \frac{f(x^0) - f^*}{f^* - f_{\text{slb}}} \leq 1 + \frac{M\text{Dist}(x^0, \text{Opt})}{f^* - f_{\text{slb}}} = 1 + MG\bar{C} \ .$$

Therefore the ratio of the new guarantee in Theorem 3.3 to the standard smoothing guarantee (25) is at most

$$\frac{\text{Guarantee of Theorem 3.3}}{\text{Standard Smoothing Guarantee (25)}} \ \leq \ \frac{8\sqrt{2}\sqrt{\varepsilon'}\ln\left(1 + MG\bar{C}\right)}{\bar{C}} \ + \ \frac{11\sqrt{2}}{\bar{C}} \ . \tag{27}$$

Notice from (26) that when $\text{Dist}(x^0, \text{Opt})$ is sufficiently large that the ratio $\bar{C}$ is very large and the right-hand side of (27) can be made arbitrarily small, thereby showing that in these cases the computational guarantee in Theorem 3.3 can be made arbitrarily better than the standard guarantee (25) for a smoothing method. In fact, even if $\bar{C}$ is not very large (say, $\bar{C} = 20$), then for very small values of ε' it will still hold that the computational guarantee in Theorem 3.3 will be better than the standard guarantee (25).

The proof of Theorem 3.3, as well as the proof of several other results to follow, will rely on the following technical lemma, whose proof is given in Appendix A.3.

Lemma 3.1. *Let $x^0 \in Q$, $B \in (0, 1)$, and relative accuracy $\varepsilon' \in (0, 1]$ be given, and consider an algorithm with outer and inner iterations (such as Algorithm 3) with initial iterate $x_{1,0} \leftarrow x^0$ and $x_{i+1,0} \leftarrow x_{i,K_i}$ where $K_i \leftarrow j$ denotes the first index j for which $\frac{f(x_{i,j}) - f_{\text{slb}}}{f(x_{i,0}) - f_{\text{slb}}} < B$. Let x^N denote the iterate value after a total of N inner iterations.*

Let A_1 and A_2 be two positive parameters satisfying $A_1 + A_2 = B$, and suppose that there are sequences $\{J_i\}$ and $\{I_i\}$ that satisfy the following conditions:

(i) for all $k \geq J_i$ it holds that $f(x_{i,k}) - f^* \leq A_1\left(f(x_{i,0}) - f_{\text{slb}}\right)$.
(ii) for all $k \geq I_i$ it holds that $f(x_{i,k}) - f^* \leq A_2\left(f(x_{i,0}) - f_{\text{slb}}\right)\varepsilon'$.

Let m denote the smallest value of i for which $K_i > J_i$ and let p denote the largest index i for which K_i exists (is finite). Then

$$m - 1 \leq \left\lfloor \frac{\ln\left(\frac{f(x_{1,0}) - f_{\text{slb}}}{f^* - f_{\text{slb}}}\right)}{\ln(1/B)} \right\rfloor \quad , \quad \textit{and} \quad p \leq m + \left\lfloor \frac{\ln\left(\frac{1}{B - A_1}\right)}{\ln(1/B)} \right\rfloor \ . \tag{28}$$

Furthermore, if $N \geq \sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i$, *it holds that* $\frac{f(x^N)-f^*}{f^*-f_{\rm slb}} \leq \varepsilon'$. □

We will prove Theorem 3.3 by using Lemma 3.1 in conjunction with the following two propositions. For notational convenience, we will work with four constants B, r_1, r_2 and t, whose specific values are $B = 4/5$ and $r_1 = r_2 = t = 1/5$.

Proposition 3.6. *Define* $T := \left\lceil G\sqrt{\frac{2A\bar{D}}{tr_1\varepsilon'}} - 1 \right\rceil$. *For* $k \geq T$ *it holds that:*

$$f(x_{i,k}) - f^* \leq (t + r_1)(f(x_{i,0}) - f_{\rm slb}) \ . \tag{29}$$

Proof: Note that for any $x \in \text{Opt}$ it holds that $f_{\mu_i}(x) \leq f^* + \bar{D}\mu_i$, whereby $\text{Opt} \subset S := \{x \in Q : f_{\mu_i}(x) \leq f^* + \bar{D}\mu_i\}$. It then follows from Theorem 2.2 applied to the function $f_{\mu_i}(\cdot)$ and using $\delta = f^* + \bar{D}\mu_i$ that for any $k \geq T$ we have:

$$\begin{aligned} f_{\mu_i}(x_{i,k}) - f^* - \bar{D}\mu_i &\leq \tfrac{2A}{\mu_i(T+1)^2}\text{Dist}(x_{i,0}, S)^2 \\ &\leq \tfrac{2A}{\mu_i(T+1)^2}\text{Dist}(x_{i,0}, \text{Opt})^2 \\ &\leq \tfrac{2A}{\mu_i(T+1)^2}G^2\left(f(x_{i,0}) - f_{\rm slb}\right)^2 \\ &= \tfrac{2A\bar{D}}{t\varepsilon'(T+1)^2}G^2\left(f(x_{i,0}) - f_{\rm slb}\right) \\ &\leq r_1\left(f(x_{i,0}) - f_{\rm slb}\right) \ , \end{aligned} \tag{30}$$

where the second inequality uses the fact that $\text{Opt} \subset S$, the third inequality uses the definition of G, the equality uses the definition of μ_i, and the last inequality uses the value of T.

Note that $\mu_i\bar{D} = t\varepsilon'\left(f(x_{i,0}) - f_{\rm slb}\right)$. Therefore

$$f(x_{i,k}) - f^* \leq f_{\mu_i}(x_{i,k}) - f^* \leq \left(t\varepsilon' + r_1\right)\left(f(x_{i,0}) - f_{\rm slb}\right) \leq (t + r_1)\left(f(x_{i,0}) - f_{\rm slb}\right) \ , \tag{31}$$

where the last inequality uses $\varepsilon' \leq 1$, completing the proof. □

Proposition 3.7. *Define* $U := \left\lceil \frac{G}{\varepsilon'}\sqrt{\frac{2A\bar{D}}{tr_2}} - 1 \right\rceil$. *For* $k \geq U$ *it holds that:*

$$f(x_{i,k}) - f^* \leq (t + r_2)(f(x_{i,0}) - f_{\rm slb})\varepsilon' \ . \tag{32}$$

Proof: Define S and δ exactly as in the start of the proof of Proposition 3.6. Using identical logic as in the proof of Proposition 3.6, it then follows from Theorem 2.2 applied to the function $f_{\mu_i}(\cdot)$

that for any $k \geq U$ we have:

$$\begin{aligned} f_{\mu_i}(x_{i,k}) - f^* - \bar{D}\mu_i &\leq \tfrac{2A}{\mu_i(U+1)^2}\mathrm{Dist}(x_{i,0}, S)^2 \\ &\leq \tfrac{2A}{\mu_i(U+1)^2}\mathrm{Dist}(x_{i,0}, \mathrm{Opt})^2 \\ &\leq \tfrac{2A}{\mu_i(U+1)^2}G^2\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)^2 \qquad (33) \\ &= \tfrac{2A\bar{D}}{t\varepsilon'(U+1)^2}G^2\left(f(x_{i,0}) - f_{\mathrm{slb}}\right) \\ &\leq r_2\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)\varepsilon' \ , \end{aligned}$$

where the second inequality uses the fact that $\mathrm{Opt} \subset S$, the third inequality uses the definition of G, the equality uses the definition of μ_i, and the last inequality uses the value of U.

Note that $\mu_i\bar{D} = t\varepsilon'\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)$. Therefore

$$f(x_{i,k}) - f^* \leq f_{\mu_i}(x_{i,k}) - f^* \leq \left(t\varepsilon' + r_2\varepsilon'\right)\left(f(x_{i,0}) - f_{\mathrm{slb}}\right) = (t + r_2)\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)\varepsilon' \ , \qquad (34)$$

completing the proof. □

Proof of Theorem 3.3: Let us invoke Lemma 3.1 applied to Algorithm 3, using $A_1 = t + r_1 = 2/5$, $A_2 = t + r_2 = 2/5$, $B = 4/5$, and $J_i = T$ and $I_i = U$ for all i. Then Propositions 3.6 and 3.7 imply that $\{J_i\}$ and $\{I_i\}$ satisfy conditions (i) and (ii) of Lemma 3.1. It remains to bound the value of $\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i$. We have:

$$\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i \ = \ (m-1)T + (p-m+2)U \ .$$

Substituting the bounds on $m-1$ and $p-m$ from (28) and the definitions of T and U from Propositions 3.6 and 3.7 into the right-hand side above yields the following upper bound:

$$\left\lfloor \frac{\ln\left(\frac{f(x_{1,0}) - f_{\mathrm{slb}}}{f^* - f_{\mathrm{slb}}}\right)}{\ln(1/B)} \right\rfloor \times \left\lceil G\sqrt{\frac{2A\bar{D}}{tr_1\varepsilon'}} - 1 \right\rceil + \left(\left\lfloor \frac{\ln\left(\frac{1}{B-A_1}\right)}{\ln(1/B)} \right\rfloor + 2 \right) \times \left\lceil \frac{G}{\varepsilon'}\sqrt{\frac{2A\bar{D}}{tr_2}} - 1 \right\rceil \ ,$$

and simplifying the numerical constants and rounding up (for numeric aesthetics) yields the bound in the theorem. □

4 Computational Guarantees when $f(\cdot)$ is Smooth

In this section we study the computational complexity of solving (1) in the case when $f(\cdot)$ is convex and differentiable on Q. We assume that $\nabla f(\cdot)$ is Lipschitz on Q as defined in (11).

Let us first consider directly applying the Accelerated Gradient Method (Algorithm 2) to solve (1), and let us apply Theorem 2.2. Let $\varepsilon' > 0$ denote the relative accuracy, and note again that an

ε'-relative solution of (1) corresponds to an absolute ε-solution for $\varepsilon := \varepsilon' \cdot (f^* - f_{\text{slb}})$. Let $x^0 \in Q$ be the initial point. It then follows from Theorem 2.2 using $\delta = f^*$ (whereby $S_\delta = \{x \in Q : f(x) \leq f^*\} = \text{Opt}$) that if

$$N \;\geq\; \frac{\sqrt{2}\sqrt{L}\text{Dist}(x^0, \text{Opt})}{\sqrt{\varepsilon'}\sqrt{f^* - f_{\text{slb}}}} \;-1\;, \tag{35}$$

then

$$\frac{f(x^N) - f^*}{f^* - f_{\text{slb}}} \;\leq\; \varepsilon'\;.$$

It turns out that we can improve the computational guarantee of (35) by using a version of the Accelerated Gradient Method that periodically restarts the method with an appropriate rule for deciding when to do the restarts. Algorithm 4 presents such a restarting method for computing an ε'-relative solution of (1) for the smooth objective function $f(\cdot)$ based on successive restarting of the Accelerated Gradient Method. At the i^{th} outer iteration of Algorithm 4 the algorithm starts the Accelerated Gradient Method at the point $x_{i,0}$ for optimizing $f(\cdot)$ on Q. The outer iteration i runs until the ratio test in Step (2.) fails, at which point the current point $x_{i,j}$ becomes the starting point of the next outer iteration, namely $x_{i+1,0} \leftarrow x_{i,j}$. The counter K_i records the number of inner iterations computed in outer iteration i. Therefore the total number of inner iterations of Algorithm 4 after completing i outer iterations is $K_1 + ... + K_i$. As was used in Algorithm 3, in the description of Algorithm 4 the notation "$x_{i,j} \leftarrow \text{AGM}(f(\cdot), x_{i,0}, j)$" denotes assigning to $x_{i,j}$ the j^{th} iterate of the Accelerated Gradient Method applied to the optimization problem (1) with objective function $f(\cdot)$ using the initial point $x_{i,0} \in Q$.

Algorithm 4 Accelerated Gradient Method with Simple Restarting

Initialize. Initialize with $x^0 \in Q$ and $\varepsilon' > 0$.
Set $x_{1,0} \leftarrow x^0$, $i \leftarrow 1$.

At outer iteration i :

1. **Initialize inner iteration.** $j \leftarrow 0$
2. **Run inner iterations.** At inner iteration j:

If $\dfrac{f(x_{i,j}) - f_{\text{slb}}}{f(x_{i,0}) - f_{\text{slb}}} \geq 0.8$, then

$x_{i,j+1} \leftarrow \text{AGM}(f(\cdot),\ x_{i,0},\ j+1)$,

$j \leftarrow j+1$, and Goto step 2.

Else $K_i \leftarrow j$, $x_{i+1,0} \leftarrow x_{i,j}$, $i \leftarrow i+1$, and Goto step 1.

Restarting for accelerated gradient methods for strongly convex functions has been studied in [17] and [23]. Our analysis works for any smooth function, and does not require the assumption of strong convexity. We have the following computational guarantee associated with Algorithm 4.

Theorem 4.1. (Complexity Bound for Accelerated Gradient Method with Simple Restarting) *Let $x^0 \in Q$ be the initial point and let x^N denote the iterate value of the Accelerated Gradient Method with Simple Restarting (Algorithm 4) after a total of N inner iterations. If*

$$N \;\geq\; G\sqrt{L}\left(17\left[\frac{\sqrt{f^* - f_{\rm slb}}}{\sqrt{\varepsilon'}}\right] \;+\; 22\sqrt{(f(x^0) - f_{\rm slb})}\right) \;,$$

then

$$\frac{f(x^N) - f^*}{f^* - f_{\rm slb}} \;\leq\; \varepsilon' \;. \qquad\qquad \square$$

Let us compare the computational guarantee of Theorem 4.1 with the standard guarantee for smoothing given in (35). Our discussion will be very similar to the discussion following Theorem 3.3. We consider cases wherein $\mathrm{Dist}(x^0, \mathrm{Opt}) > G \cdot (f^* - f_{\rm slb})$. Since G and $f^* - f_{\rm slb}$ are themselves fixed for a given problem instance, this means that $\mathrm{Dist}(x^0, \mathrm{Opt})$ is large, i.e., x^0 is far from the set of optimal solutions. Let $\bar{C}$ denote the ratio of $\mathrm{Dist}(x^0, \mathrm{Opt})$ to $G \cdot (f^* - f_{\rm slb})$, namely

$$\bar{C} := \frac{\mathrm{Dist}(x^0, \mathrm{Opt})}{G(f^* - f_{\rm slb})} > 1 \;. \tag{36}$$

Note that $f(x^0) \leq f^* + \frac{L}{2}\mathrm{Dist}(x^0, \mathrm{Opt})^2$, and therefore

$$\frac{f(x^0) - f_{\rm slb}}{f^* - f_{\rm slb}} \leq 1 + \tfrac{L}{2}\bar{C}^2 G^2 (f^* - f_{\rm slb}) \;.$$

Therefore the ratio of the new guarantee in Theorem 4.1 to the standard guarantee for the Accelerated Gradient Method (35) is at most

$$\begin{aligned}
\frac{\text{Guarantee of Theorem 4.1}}{\text{Standard AGM Guarantee (35)}} \;&=\; \frac{8.5\sqrt{2}}{\bar{C}} + 11\sqrt{\varepsilon'}\sqrt{L}G\sqrt{f^* - f_{\rm slb}}\sqrt{1 + \frac{1}{\frac{L}{2}\bar{C}^2 G^2 (f^* - f_{\rm slb})}} \\
&\leq\; \frac{8.5\sqrt{2}}{\bar{C}} + 11\sqrt{\varepsilon'}\sqrt{L}G\sqrt{f^* - f_{\rm slb}}\left(1 + \frac{1}{L\bar{C}^2 G^2 (f^* - f_{\rm slb})}\right) \\
&=\; \frac{8.5\sqrt{2}}{\bar{C}} + 11\sqrt{\varepsilon'}\sqrt{L}G\sqrt{f^* - f_{\rm slb}} \;+\; \frac{11\sqrt{\varepsilon'}}{\bar{C}^2 G\sqrt{L}\sqrt{f^* - f_{\rm slb}}} \;,
\end{aligned} \tag{37}$$

where the inequality uses $\sqrt{1 + \frac{1}{t}} \leq 1 + \frac{1}{2t}$. Let us consider cases when ε' is small and $\mathrm{Dist}(x^0, \mathrm{Opt})$ is sufficiently large that the ratio $\bar{C}$ from (36) is very large. Then the first and third terms in the right-hand side of (37) can be made arbitrarily small due to a large value of $\bar{C}$, and the second term can be made arbitrarily small due to a small value of ε', thereby showing that in these cases the computational guarantee in Theorem 4.1 can be made arbitrarily better than the standard guarantee (35) for the Accelerated Gradient Method. In fact, even if $\bar{C}$ is not very large (say, $\bar{C} = 20$), then for very small values of ε' it will still hold that the computational guarantee in Theorem 4.1 will be better than the standard guarantee (35).

We will prove Theorem 4.1 by using Lemma 3.1 in conjunction with the following two propositions. For notational convenience, we will work with three constants B, r_1, and r_2, whose specific values are $B = 4/5$ and $r_1 = r_2 = 2/5$.

Proposition 4.1. *Define* $J_i := \left\lceil G\sqrt{\frac{2L(f(x_{i,0})-f_{\rm slb})}{r_1}} - 1 \right\rceil$. *For* $k \geq J_i$, *it holds that:*

$$f(x_{i,k}) - f^* \leq r_1(f(x_{i,0}) - f_{\rm slb}) \ . \tag{38}$$

Proof: It follows from Theorem 2.2 applied to the function $f(\cdot)$ and using $\delta = f^*$ that for any $k \geq J_i$ we have:

$$f(x_{i,k}) - f^* \leq \frac{2L}{(J_i+1)^2}\mathrm{Dist}(x_{i,0},\mathrm{Opt})^2 \leq \frac{2L}{(J_i+1)^2}G^2\left(f(x_{i,0}) - f_{\rm slb}\right)^2 \leq r_1\left(f(x_{i,0}) - f_{\rm slb}\right) \ , \tag{39}$$

where the second inequality uses the definition of G, and the last inequality uses the value of J_i. □

Proposition 4.2. *Define* $I_i := \left\lceil G\sqrt{\frac{2L(f(x_{i,0})-f_{\rm slb})}{r_2\varepsilon'}} - 1 \right\rceil$. *For* $k \geq I_i$ *it holds that:*

$$f(x_{i,k}) - f^* \leq r_2(f(x_{i,0}) - f_{\rm slb})\varepsilon' \ . \tag{40}$$

Proof: This proof is similar to the proof of Proposition 4.1. For any $k \geq I_i$ we have:

$$f(x_{i,k}) - f^* \leq \frac{2L}{(I_i+1)^2}\mathrm{Dist}(x_{i,0},\mathrm{Opt})^2 \leq \frac{2L}{(I_i+1)^2}G^2\left(f(x_{i,0}) - f_{\rm slb}\right)^2 \leq r_2\left(f(x_{i,0}) - f_{\rm slb}\right)\varepsilon' \ , \tag{41}$$

where the last inequality uses the value of I_i . □

Proof of Theorem 4.1: Let us invoke Lemma 3.1 applied to Algorithm 4, using $A_1 = r_1 = 2/5$, $A_2 = r_2 = 2/5$, $B = 4/5$, and $\{J_i\}$ and $\{I_i\}$ as defined in Propositions 4.1 and 4.2. Note that these two propositions imply that $\{J_i\}$ and $\{I_i\}$ satisfy conditions (i) and (ii) of Lemma 3.1. It therefore remains to bound the value of $\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i$. We have:

$$\begin{aligned}
\sum_{i=1}^{m-1} J_i &\leq \sum_{i=1}^{m-1} G\sqrt{\frac{2L(f(x_{i,0}) - f_{\rm slb})}{r_1}} \\
&\leq \left(\sum_{i=1}^{m-1} B^{\frac{i-1}{2}}\right) G\sqrt{\tfrac{2L(f(x_{1,0})-f_{\rm slb})}{r_1}} \\
&< \left(\sum_{i=1}^{\infty} B^{\frac{i-1}{2}}\right) G\sqrt{\tfrac{2L(f(x_{1,0})-f_{\rm slb})}{r_1}} \quad = \quad \tfrac{G}{1-\sqrt{B}}\sqrt{\tfrac{2L(f(x_{1,0})-f_{\rm slb})}{r_1}} \ ,
\end{aligned} \tag{42}$$

where the first inequality uses the definition of J_i and the second inequality is due to the fact that $f(x_{i,0}) - f_{\rm slb} \leq B^{i-1}(f(x_{1,0}) - f_{\rm slb})$. We also have:

$$\begin{aligned}
\sum_{i=m}^{p+1} I_i &\leq \sum_{i=m}^{p+1} G\sqrt{\frac{2L(f(x_{i,0}) - f_{\rm slb})}{r_2\varepsilon'}} \\
&\leq \left(\sum_{i=0}^{p+1-m} B^{\frac{i}{2}}\right) G\sqrt{\tfrac{2L(f(x_{m,0})-f_{\rm slb})}{r_2\varepsilon'}} \quad \leq \quad \left(\sum_{i=0}^{p+1-m} B^{\frac{i}{2}}\right) G\sqrt{\tfrac{2L(f^*-f_{\rm slb})}{(B-r_1)r_2\varepsilon'}} \ ,
\end{aligned} \tag{43}$$

where the first inequality uses the definition of I_i, the second inequality is from the fact that $f(x_{i+1,0}) - f_{\text{slb}} \leq B(f(x_{i,0}) - f_{\text{slb}})$, and the third inequality uses (60) from the proof of Lemma 3.1. Next notice from the values of the constants $B = 4/5$ and $A_1 = 2/5$ that it follows from (28) of Lemma 3.1 that $p - m + 1 \leq 5$. Therefore we have from (42) and (43):

$$\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i \;\; \leq \;\; \frac{G}{1-\sqrt{B}}\sqrt{\frac{2L(f(x_{1,0}) - f_{\text{slb}})}{r_1}} + \left(\sum_{i=0}^{5} B^{\frac{i}{2}}\right) G\sqrt{\frac{2L(f^* - f_{\text{slb}})}{(B-r_1)r_2\varepsilon'}} \;\; ,$$

and simplifying the numerical constants and rounding up (for numeric aesthetics) yields the bound stated in the theorem. □

It turns out that we can further improve the computational guarantee of Theorem 4.1 by further modifying the Accelerated Gradient Method with Simple Restarting (Algorithm 4), if we know and can easily work with an adjoint representation of $f(\cdot)$. Let us see how this can be done. We will assume that $f(\cdot)$ has the representation:

$$f(x) = \max_{\lambda \in P}\{\lambda^T Ax - d(\lambda)\} \;\; , \tag{44}$$

where P is a convex set and $d(\cdot)$ is a strongly convex function on P with strong convexity parameter σ and for which $\min_{\lambda \in P} d(\lambda) \geq 0$. (See [14] for properties of strongly convex functions.) It then follows that $f(\cdot)$ is a globally smooth convex function with Lipschitz constant at most $L := \|A\|^2/\sigma$, see Nesterov [15]. We presume further that A, $d(\cdot)$, and P are given and that the optimization problem in (44) is simple to solve. That being the case, for a given $x \in Q$, if $\tilde{\lambda}$ solves the optimization problem (44), then it holds that $f(x) = \tilde{\lambda}^T Ax - d(\tilde{\lambda})$ and $\nabla f(x) = A^T\tilde{\lambda}$.

In a similar spirit as the smoothing technique employed in Section 3.2, we will consider parametrically working with a modification $f_\mu(\cdot)$ of $f(\cdot)$ that is more smooth than $f(\cdot)$ by increasing the weight on the the strongly convex function $d(\cdot)$ in (44). For any $\mu \geq 0$ define the function $f_\mu(\cdot)$ by:

$$f_\mu(x) = \max_{\lambda \in P}\{\lambda^T Ax - (1+\mu)d(\lambda)\} \;\; . \tag{45}$$

Then $f(\cdot) = f_0(\cdot)$, and $f_\mu(\cdot)$ has a Lipschitz gradient with constant at most $L_\mu := L/(1+\mu)$, and hence is more smooth than $f(\cdot)$. Indeed, the Lipschitz constant L_μ is strictly decreasing in μ for $\mu \geq 0$, and so $f_\mu(\cdot)$ is more smooth for larger values of μ. Furthermore, if P is bounded, then $\bar{D} := \max_{\lambda \in P}\{d(\lambda)\}$ is finite, and it easily follows that

$$f(x) - \mu\bar{D} \leq f_\mu(x) \leq f(x) \text{ for all } x \;\; . \tag{46}$$

[Note that even if the representation (44) is not explicitly known, it is guaranteed to exist if $f(\cdot)$ is globally smooth. Indeed, if $f(\cdot)$ is a globally smooth convex function with Lipschitz gradient with constant L, then $d(\lambda) := \max_x\{\lambda^T Ix - f(x)\}$ is $\sigma := 1/L$-strongly convex and it also follows in this case that $f(x) = \max_{\lambda \in P}\{\lambda^T Ix - d(\lambda)\}$ where $P := \{\nabla f(x) : x \in E\}$ and E is the Euclidean space wherein x resides. For a lucid exposition of these notions, see [12].]

Let $\varepsilon' > 0$ be given. As before, we aspire to compute an ε'-relative solution of (1) as defined in (2). Algorithm 5 presents a method for computing an ε'-relative solution of (1) for the smooth objective function $f(\cdot)$ based on successive parametric smoothing and re-starting of the Accelerated Gradient Method (Algorithm 2) that is very similar in spirit to Algorithm 3. At the i^{th} outer iteration of Algorithm 5 the algorithm sets the smoothing parameter μ_i in Step (1.) and then starts the Accelerated Gradient Method at the point $x_{i,0}$ for optimizing $f_{\mu_i}(\cdot)$ on Q. Notice that the smoothing parameter μ_i decreases over the course of the outer iterations, as it makes more sense to set μ_i higher at first and then decrease it as the solution is approached. The outer iteration i runs until the ratio test in Step (3.) fails, at which point the current point $x_{i,j}$ becomes the starting point of the next outer iteration, namely $x_{i+1,0} \leftarrow x_{i,j}$. The counter K_i records the number of inner iterations computed in outer iteration i. Therefore the total number of inner iterations of Algorithm 5 after compleing i outer iterations is $K_1 + \ldots + K_i$.

As in previous descriptions, in Algorithm 5 the notation "$x_{i,j} \leftarrow \text{AGM}(f_{\mu_i}(\cdot),\ x_{i,0},\ j)$" denotes assigning to $x_{i,j}$ the j^{th} iterate of the Accelerated Gradient Method applied to the optimization problem (1) with objective function $f_{\mu_i}(\cdot)$ using the initial point $x^0 \in Q$.

Algorithm 5 Scheme for Smooth Optimization via Parametric Increased Smoothing/Restarting

---

Initialize. Initialize with $x^0 \in Q$ and $\varepsilon' > 0$.
Set $x_{1,0} \leftarrow x^0$, $i \leftarrow 1$.

At outer iteration i :

1. **Set smoothing parameter.** $\mu_i \leftarrow \dfrac{\varepsilon' \cdot (f(x_{i,0}) - f_{\text{slb}})}{5\bar{D}}$.
2. **Initialize inner iteration.** $j \leftarrow 0$
3. **Run inner iterations.** At inner iteration j:

   If $\dfrac{f(x_{i,j}) - f_{\text{slb}}}{f(x_{i,0}) - f_{\text{slb}}} \geq 0.8$, then

   $x_{i,j+1} \leftarrow \text{AGM}(f_{\mu_i}(\cdot),\ x_{i,0},\ j+1)$,

   $j \leftarrow j+1$, and Goto step 3.

   Else $K_i \leftarrow j$, $x_{i+1,0} \leftarrow x_{i,j}$, $i \leftarrow i+1$, and Goto step 1.

We have the following computational guarantee associated with Algorithm 5.

Theorem 4.2. (Complexity Bound for Scheme for Smooth Optimization via Parametric Increased Smoothing/Restarting) *Let $x^0 \in Q$ be the initial point and let the relative accuracy $\varepsilon' \in (0,1]$ be given, and let x^N denote the iterate value of the Scheme for Smooth Optimization (Algorithm 5) after a total of N inner iterations. If*

$$N \ \geq \ G\sqrt{L}\left(24\left[\frac{\sqrt{f^* - f_{\text{slb}}}}{\sqrt{\varepsilon'}}\right] + 32\left[\frac{\sqrt{\bar{D}}\ln\left(1 + \frac{f(x^0)-f^*}{f^*-f_{\text{slb}}}\right)}{\sqrt{\varepsilon'}}\right]\right) \ ,$$

then

$$\frac{f(x^N) - f^*}{f^* - f_{\text{slb}}} \ \leq \ \varepsilon' \ . \qquad \square$$

Let us now briefly compare the computational guarantee for the Accelerated Gradient Method with Parametric Smoothing/Restarting (Theorem 4.2), with the guarantee for the Accelerated Gradient Method with Simple Restarting (Theorem 4.1). Note that up to a constant factor the guarantee in Theorem 4.2 differs from that of Theorem 4.1 only in the second term. In cases where $f(x^0) - f^*$ is very large and ε' is not very small, it follows that the second term of the bound in Theorem 4.2 will be significantly smaller than that of Theorem 4.1. In these cases the Accelerated Gradient Method with Parametric Smoothing/Restarting (Theorem 4.2) yields an improved computational guarantee compared to the Accelerated Gradient Method with Simple Restarting (Theorem 4.1).

We will prove Theorem 4.2 by using Lemma 3.1 in conjunction with the following two propositions. Very similar to the proof of Theorem 3.3, for notational convenience we will work with four constants B, r_1, r_2 and t, whose specific values are $B = 4/5$ and $r_1 = r_2 = t = 1/5$.

Proposition 4.3. *Define* $T := \left\lceil G\sqrt{\frac{2L\bar{D}}{tr_1\varepsilon'}} - 1 \right\rceil$. *For* $k \geq T$ *it holds that:*

$$f(x_{i,k}) - f^* \leq (t + r_1)(f(x_{i,0}) - f_{\text{slb}}) \ . \tag{47}$$

Proof: Note that for any $x \in \text{Opt}$ it holds that $f_{\mu_i}(x) \leq f^*$, whereby $\text{Opt} \subset S := \{x \in Q : f_{\mu_i}(x) \leq f^*\}$. It follows from Theorem 2.2 applied to the function $f_{\mu_i}(\cdot)$ and using $\delta = f^*$ that for any $k \geq T$ we have:

$$\begin{aligned} f_{\mu_i}(x_{i,k}) - f^* &\leq \tfrac{2L}{(1+\mu_i)(T+1)^2}\text{Dist}(x_{i,0}, S)^2 \\ &\leq \tfrac{2L}{(1+\mu_i)(T+1)^2}\text{Dist}(x_{i,0}, \text{Opt})^2 \\ &< \tfrac{2L}{\mu_i(T+1)^2}G^2\left(f(x_{i,0}) - f_{\text{slb}}\right)^2 \\ &= \tfrac{2L\bar{D}}{t\varepsilon'(T+1)^2}G^2\left(f(x_{i,0}) - f_{\text{slb}}\right) \\ &\leq r_1\left(f(x_{i,0}) - f_{\text{slb}}\right) \ , \end{aligned} \tag{48}$$

where the second inequality uses the fact that $\text{Opt} \subset S$, the third inequality uses the definition of G and $1 + \mu_i > \mu_i$, the equality uses the definition of μ_i, and the last inequality uses the value of T. Recall that $f(x) - \mu_i\bar{D} \leq f_{\mu_i}(x)$. Therefore for any $k \geq T$ we have:

$$f(x_{i,k}) - f^* \leq f_{\mu_i}(x_{i,k}) - f^* + \mu_i\bar{D} \leq \left(t\varepsilon' + r_1\right)\left(f(x_{i,0}) - f_{\text{slb}}\right) \leq (t + r_1)\left(f(x_{i,0}) - f_{\text{slb}}\right) \ , \tag{49}$$

where the last inequality uses $\varepsilon' \leq 1$, completing the proof. □

Proposition 4.4. *Define* $I_i := \left\lceil G\sqrt{\frac{2L(f(x_{i,0}) - f_{\text{slb}})}{r_2\varepsilon'}} - 1 \right\rceil$. *For* $k \geq I_i$ *it holds that:*

$$f(x_{i,k}) - f^* \leq (t + r_2)(f(x_{i,0}) - f_{\text{slb}})\varepsilon' \ . \tag{50}$$

This proof is similar to the proof of Proposition 4.3. For any $k \geq I_i$ we have:

$$\begin{aligned} f_{\mu_i}(x_{i,k}) - f^* &\leq \tfrac{2L}{(1+\mu_i)(I_i+1)^2}\mathrm{Dist}(x_{i,0}, S)^2 \\ &\leq \tfrac{2L}{(1+\mu_i)(I_i+1)^2}\mathrm{Dist}(x_{i,0}, \mathrm{Opt})^2 \\ &< \tfrac{2L}{(I_i+1)^2}G^2\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)^2 \\ &\leq r_2\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)\varepsilon' \ , \end{aligned} \tag{51}$$

where the last inequality uses the value of I_i . Therefore:

$$f(x_{i,k}) - f^* \leq f_{\mu_i}(x_{i,k}) - f^* + \mu_i\bar{D} \leq \left(t\varepsilon' + r_2\varepsilon'\right)\left(f(x_{i,0}) - f_{\mathrm{slb}}\right) = (t + r_2)\left(f(x_{i,0}) - f_{\mathrm{slb}}\right)\varepsilon' \ , \tag{52}$$

which completes the proof. □

Proof of Theorem 4.2: Let us invoke Lemma 3.1 applied to Algorithm 5, using $A_1 = r_1 + t = 2/5$, $A_2 = r_2 + t = 2/5$, $B = 4/5$, and $\{J_i\}$ defined by $J_i = T$ for all i, and $\{I_i\}$ as defined in Propositions 4.3 and 4.4. Note that these two propositions imply that $\{J_i\}$ and $\{I_i\}$ satisfy conditions (i) and (ii) of Lemma 3.1. It therefore remains to bound the value of $\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i$. We have from (28) of Lemma 3.1 that:

$$\sum_{i=1}^{m-1} J_i = (m-1)T \leq \left\lfloor \frac{\ln\left(\dfrac{f(x_{1,0}) - f_{\mathrm{slb}}}{f^* - f_{\mathrm{slb}}}\right)}{\ln(1/B)} \right\rfloor \times \left\lceil G\sqrt{\frac{2L\bar{D}}{tr_1\varepsilon'}} - 1 \right\rceil \ , \tag{53}$$

and also

$$\begin{aligned} \sum_{i=m}^{p+1} I_i &\leq \sum_{i=m}^{p+1} G\sqrt{\frac{2L(f(x_{i,0}) - f_{\mathrm{slb}})}{r_2\varepsilon'}} \\ &\leq \left(\sum_{i=0}^{p+1-m} B^{\frac{i}{2}}\right) G\sqrt{\tfrac{2L(f(x_{m,0}) - f_{\mathrm{slb}})}{r_2\varepsilon'}} \\ &\leq \left(\sum_{i=0}^{p+1-m} B^{\frac{i}{2}}\right) G\sqrt{\tfrac{2L(f^* - f_{\mathrm{slb}})}{(B-t-r_1)r_2\varepsilon'}} \ , \end{aligned} \tag{54}$$

where the first inequality uses the definition of I_i, the second inequality is from the fact that $f(x_{i+1,0}) - f_{\mathrm{slb}} \leq B(f(x_{i,0}) - f_{\mathrm{slb}})$, and the third inequality uses (60) from the proof of Lemma 3.1.

Next notice from the values of the constants $B = 4/5$ and $A_1 = 2/5$ that it follows from (28) of Lemma 3.1 that $p - m + 1 \leq 5$. Therefore we have from (53) and (54):

$$\sum_{i=1}^{m-1} J_i + \sum_{i=m}^{p+1} I_i \leq \left\lfloor \frac{\ln\left(\frac{f(x_{1,0}) - f_{\mathrm{slb}}}{f^* - f_{\mathrm{slb}}}\right)}{\ln(1/B)} \right\rfloor \times \left\lceil G\sqrt{\tfrac{2L\bar{D}}{tr_1\varepsilon'}} - 1 \right\rceil + \left(\sum_{i=0}^{5} B^{\frac{i}{2}}\right) G\sqrt{\tfrac{2L(f^* - f_{\mathrm{slb}})}{(B-t-r_1)r_2\varepsilon'}} \ ,$$

and simplifying the numerical constants and rounding up (for numeric aesthetics) yields the bound stated in the theorem. □

A Appendix

A.1 Growth Constant G in Relation to the Modulus of Weak Sharp Minima

The optimal solution set Opt of (1) is called a set of *weak sharp minima* with modulus α if it holds that:

$$f(x) \geq f^* + \alpha \cdot \text{Dist}(x, \text{Opt}) \quad \text{for all } x \in Q \ . \tag{55}$$

This concept was first developed by Polyak [18] when Opt is a singleton, and generalized by Burke and Ferris [3] to include the possibility of multiple optima. The modulus of weak sharp minima has been a useful tool in sensitivity analysis [8, 11], convergence analysis for certain problem classes [7, 3], linear regularity and error bounds [4, 5, 6], perturbation properties of nonlinear optimization [21, 22, 2], as well as in the finite termination of certain algorithms [19], [9], and [7].

Comparing (55) to (5), we see that the modulus α of weak sharp minima is a close cousin of the growth constant G. Indeed, if we were to loosen the restriction that f_{slb} be a strict lower bound and instead allow it to take the value $f_{\text{slb}} = f^*$ in the definition of G in (3), then we would obtain precisely that $G = \alpha^{-1}$. However, the notion of f_{slb} being a strict lower bound is fundamental for the results herein.

Note that (55) specifies the exact local growth of $f(\cdot)$ away from the set of optimal solutions. And although as defined in (55) the weak sharp minima is a global property, due to convexity it is essentially a local property and indeed its usefulness derives from the local nature of the weak sharp minima in a neighborhood of the optimal solution set. This is in contrast to the growth constant G as defined in (3), which by its nature is a global property as illustrated in the constructions in Figure 1. Last of all we point out that while one can easily have $\alpha = 0$ for weak sharp minima (just let $f(x)$ be a differentiable convex function whose optimum is attained in the relative interior of Q), Theorem 1.1 shows that G is finite for all reasonably-behaved convex functions.

A.2 Proof of Theorem 1.1

Proof of Theorem 1.1: Let us fix an optimal solution $x^* \in \text{Opt}$, and define $\delta := \max_{v \in E_\varepsilon}\{\|v - x^*\|\}$ and define $\bar{G} := \max\{\frac{\delta}{\varepsilon}, \frac{\delta}{f^* - f_{\text{slb}}}\}$. We will prove that for any $x \in Q$, the following inequality holds:

$$\text{Dist}\,(x, \text{Opt}) \leq \bar{G}\,(f(x) - f_{\text{slb}}) \ , \tag{56}$$

which then implies that $G \leq \bar{G}$ is finite. We consider two cases as follows:

Case (i): $x \in \text{Opt}_\varepsilon$. In this case we have $x = v + s$ where $v \in E_\varepsilon$ and $s \in S$. Since s is in the recession cone of Opt_ε it holds that $x^* + s \in \text{Opt}$, whereby

$$\text{Dist}(x, \text{Opt}) \leq \|x - (x^* + s)\| = \|v - x^*\| \leq \delta \ , \tag{57}$$

and therefore

$$f(x) - f_{\text{slb}} \geq f^* - f_{\text{slb}} \geq \frac{(f^* - f_{\text{slb}})(\text{Dist}(x, \text{Opt}))}{\delta} \geq \bar{G}^{-1}\text{Dist}(x, \text{Opt}) \ ,$$

which shows (56) in this case.

Case (ii): $x \notin \text{Opt}_\varepsilon$. Let x^1 be the projection of x onto Opt and let x^2 be the point on the line segment from x^1 to x that satisfies $f(x^2) = f^* + \varepsilon$. (Existence of x^2 is guaranteed by continuity of $f(\cdot)$.) Then

$$f(x) - f_{\text{slb}} \geq f(x) - f^* \geq (f(x^2) - f^*)\frac{\|x - x^1\|}{\|x^2 - x^1\|} \geq \frac{\varepsilon\|x - x^1\|}{\delta} = \frac{\varepsilon\text{Dist}(x, \text{Opt})}{\delta} \geq \bar{G}^{-1}\text{Dist}(x, \text{Opt}) \ ,$$

where the second inequality is from the convexity of $f(\cdot)$ which implies the chordal inequality $\frac{f(x)-f^*}{\|x-x^1\|} \geq \frac{f(x^2)-f^*}{\|x^2-x^1\|}$, and the third inequality uses $\|x^2 - x^1\| = \text{Dist}(x^2, \text{Opt}) \leq \delta$ (from (57)). The last equality above uses the fact that $\text{Dist}(x, \text{Opt}) = \|x - x^1\|$. This proves (56) in this case. □

A.3 Proof of Lemma 3.1

We first prove the following:

Proposition A.1. *If* $\frac{f(x_{i,0})-f_{\text{slb}}}{f^*-f_{\text{slb}}} \leq \frac{1}{B-A_1}$, *then either* $\frac{f(x_{i,I_i})-f^*}{f^*-f_{\text{slb}}} \leq \varepsilon'$ *or* $K_i \leq I_i$.

Proof: Suppose that $K_i > I_i$. Thus

$$f(x_{i,I_i}) - f^* \leq A_2(f(x_{i,0}) - f_{\text{slb}})\varepsilon' \leq \frac{A_2}{B - A_1}\,(f^* - f_{\text{slb}})\,\varepsilon' = (f^* - f_{\text{slb}})\,\varepsilon' \ ,$$

where we use the fact that $B = A_1 + A_2$. □

Proof of Lemma 3.1: It follows using induction that $f(x_{i,0}) - f_{\text{slb}} \leq B^{i-1}(f(x_{1,0}) - f_{\text{slb}})$. Therefore $f^* - f_{\text{slb}} \leq f(x_{m,0}) - f_{\text{slb}} \leq B^{m-1}(f(x_{1,0}) - f_{\text{slb}})$, and taking logarithms yields

$$m - 1 \leq \frac{\ln\left(\dfrac{f(x_{1,0}) - f_{\text{slb}}}{f^* - f_{\text{slb}}}\right)}{\ln(1/B)} \ ,$$

which proves the first inequality of (28). Next note that $K_m - 1 \geq J_m$, and therefore

$$f(x_{m,K_m-1}) - f^* \leq A_1\,(f(x_{m,0}) - f_{\text{slb}}) \ , \tag{58}$$

and also

$$\frac{f(x_{m,K_m-1}) - f_{\text{slb}}}{f(x_{m,0}) - f_{\text{slb}}} \geq B \ . \tag{59}$$

It then follows that

$$\frac{f(x_{m,0}) - f_{\text{slb}}}{f^* - f_{\text{slb}}} = \frac{1}{\frac{f(x_{m,K_m-1})-f_{\text{slb}}}{f(x_{m,0})-f_{\text{slb}}} - \frac{f(x_{m,K_m-1})-f^*}{f(x_{m,0})-f_{\text{slb}}}} \leq \frac{1}{B - A_1} \ , \tag{60}$$

where the inequality uses (59) and (58).

Let us now prove the second inequality of (28). If $p < m$ the second inequality of (28) follows trivially. If $p \geq m$ we have $f^* - f_{\text{slb}} \leq f(x_{p,0}) - f_{\text{slb}} \leq B^{p-m}(f(x_{m,0}) - f_{\text{slb}}) \leq B^{p-m}\frac{1}{B-A_1}(f^* - f_{\text{slb}})$ from (60). Taking logarithms yields

$$p - m \leq \frac{\ln\left(\frac{1}{B-A_1}\right)}{\ln(1/B)} ,$$

which proves the second inequality of (28).

Define $N_1 := \sum_{i=1}^{m-1} K_i$. It follows from the definition of m that $K_i \leq J_i$ for $i = 1, \ldots, m-1$, and hence

$$N_1 \leq \sum_{i=1}^{m-1} J_i =: \hat{N}_1 \tag{61}$$

where $\hat{N}_1$ is defined to be the term on the right above. Define $N_2 := \sum_{i=m}^{p+1} I_i$. From (60) we know that $\frac{f(x_{m,0})-f_{\text{slb}}}{f^*-f_{\text{slb}}} \leq \frac{1}{B-A_1}$, and hence for $m \leq i \leq p+1$ it also holds that $\frac{f(x_{i,0})-f_{\text{slb}}}{f^*-f_{\text{slb}}} \leq \frac{1}{B-A_1}$. Therefore it follows from Proposition A.1 that by the time the algorithm has computed a total of $N_1 + N_2$ inner iterations that $\frac{f(x_{i,I_i})-f^*}{f^*-f_{\text{slb}}} \leq \varepsilon'$ for some $i \leq p+1$. It therefore suffices to bound $N_1 + N_2$. We have $N_1 + N_2 \leq \hat{N}_1 + N_2$ from (61). □

References

[1] Amir Beck and Marc Teboulle, *Smoothing and first order methods: A unified framework*, SIAM Journal on Optimization **22** (2012), no. 2, 557–580.

[2] J. Frederic Bonnans and Alexander D. Ioffe, *Quadratic growth and stability in convex programming problems with multiple solutions*, Tech. report, INRIA Research Report RR-2403, 1994.

[3] J. V. Burke and M.C. Ferris, *Weak sharp minima in mathematical programming*, SIAM Journal on Control and Optimization **31** (1993), no. 5, 1340–1359.

[4] James V. Burke and Sien Deng, *Weak sharp minima revisited, part I: basic theory*, Control and Cybernetics **31** (2002), no. 3, 439–469.

[5] ———, *Weak sharp minima revisited, part II: application to linear regularity and error bounds*, Mathematical programming **104** (2005), no. 2-3, 235–261.

[6] ———, *Weak sharp minima revisited, part III: Error bounds for differentiable convex inclusions*, Mathematical Programming **116** (2009), no. 1-2, 37–56.

[7] James V. Burke and Michael C. Ferris, *A Gauss-Newton method for convex composite optimization*, Mathematical Programming **71** (1995), no. 2, 179–194.

[8] James V. Burke, Adrian S. Lewis, and Michael L. Overton, *Optimal stability and eigenvalue multiplicity*, Foundations of Computational Mathematics **1** (2001), no. 2, 205–225.

[9] Michael C. Ferris, *Finite termination of the proximal point algorithm*, Mathematical Programming **50** (1991), no. 1-3, 359–366.

[10] Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *Elements of statistical learning*, second ed., Springer Series in Statistics, Springer, New York, 2009.

[11] Abderrahim Jourani, *Hoffman's error bound, local controllability, and sensitivity analysis*, SIAM Journal on Control and Optimization **38** (2000), no. 3, 947–970.

[12] Sham M. Kakade, Shai Shalev-Shwartz, and Ambuj Tewari, *Regularization techniques for learning with matrices*, The Journal of Machine Learning Research **13** (2012), no. 1, 1865–1890.

[13] A. S. Nemirovsky and D. B. Yudin, *Problem complexity and method efficiency in optimization*, Wiley, New York, 1983.

[14] Y.E. Nesterov, *Introductory lectures on convex optimization: a basic course*, Kluwer Academic Publishers, Boston, 2003.

[15] Y.E. Nesterov, *Smooth minimization of non-smooth functions*, Mathematical Programming **103** (2005), no. 1, 127–152.

[16] Yurii Nesterov, *Smoothing technique and its applications in semidefinite optimization*, Mathematical Programming **110** (2007), no. 2, 245–259.

[17] Brendan O'Donoghue and Emmanuel Candes, *Adaptive restart for accelerated gradient schemes*, Foundations of Computational Mathematics **15** (2013), no. 3, 715–732.

[18] B. Polyak, *Sharp minima*, Proceedings of the IIASA Workshop on Generalized Lagrangians and Their Applications, Laxenburg, Austria. Institute of Control Sciences Lecture Notes, Moscow, 1979.

[19] B. Polyak, *Introduction to optimization*, Optimization Software, Inc., New York, 1987.

[20] J. Renegar, *A framework for applying subgradient methods to conic optimization problems, version 2*, Tech. report, Cornell University, 2015.

[21] Alexander Shapiro, *Perturbation theory of nonlinear programs when the set of optimal solutions is not a singleton*, Applied Mathematics and Optimization **18** (1988), no. 1, 215–229.

[22] ———, *Perturbation analysis of optimization problems in Banach spaces*, Numerical Functional Analysis and Optimization **13** (1992), no. 1-2, 97–116.

[23] Weijie Su, Stephen Boyd, and Emmanuel Candes, *A differential equation for modeling Nesterov's accelerated gradient method: theory and insights*, Advances in Neural Information Processing Systems, 2014, pp. 2510–2518.

[24] P. Tseng, *On accelerated proximal gradient methods for convex-concave optimization*, Tech. report, May 21, 2008.